



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/147108/>

Version: Published Version

Article:

Connor, Stephen Bryan and Pymar, Richard (2019) Mixing times for exclusion processes on hypergraphs. *Electronic Journal of Probability*. 73. ISSN: 1083-6489

<https://doi.org/10.1214/19-EJP332>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Mixing times for exclusion processes on hypergraphs*

Stephen B. Connor[†] Richard J. Pymar[‡]

Abstract

We introduce a natural extension of the exclusion process to hypergraphs and prove an upper bound for its mixing time. In particular we show the existence of a constant C such that for any connected, regular hypergraph G within some natural class, the ε -mixing time of the exclusion process on G with any feasible number of particles can be upper-bounded by $CT_{\text{EX}(2,G)} \log(|V|/\varepsilon)$, where $|V|$ is the number of vertices in G and $T_{\text{EX}(2,G)}$ is the $1/4$ -mixing time of the corresponding exclusion process with just two particles. Moreover we show this is optimal in the sense that there exist hypergraphs in the same class for which $T_{\text{EX}(2,G)}$ and the mixing time of just one particle are not comparable. The proofs involve an adaptation of the *chameleon process*, a technical tool invented by Morris ([14]) and developed by Oliveira ([15]) for studying the exclusion process on a graph.

Keywords: mixing time; exclusion; interchange; random walk; hypergraph; coupling.

AMS MSC 2010: Primary 60J27; 60K35, Secondary 82C22.

Submitted to EJP on March 23, 2018, final version accepted on June 9, 2019.

Supersedes arXiv:1606.02703v4.

1 Introduction

Let $G = (V, E)$ be a finite connected graph with vertex set V and edge set E . Fix $k \in \{1, \dots, |V|\}$ and consider k indistinguishable particles moving on V using the following rules:

1. each vertex is occupied by at most one particle,
2. each edge $e \in E$ rings at the times of a Poisson process of rate 1, independently of all other edges,

*Both authors supported in part by LMS Research in Pairs (Scheme 4) Grant 41215. SBC supported in part by EPSRC Research Grant EP/J009180. RJP supported in part by Leverhulme Research Grant RPG-2012-608 held by Nadia Sidorova and EPSRC Grant EP/M027694/1 held by Codina Cotar.

[†]Department of Mathematics, University of York, York, YO10 5DD, UK. E-mail: stephen.connor@york.ac.uk

[‡]Department of Economics, Mathematics and Statistics, Birkbeck, University of London, London, WC1E 7HX, UK. E-mail: r.pymar@bbk.ac.uk

3. when an edge $e = \{u, v\}$ rings, the occupancy states of vertices u and v are switched.

For each $v \in V$ and $t \geq 0$, let $\eta_t(v) = 1$ if v is occupied at time t , and $\eta_t(v) = 0$ if v is vacant at time t . The process $(\eta_t)_{t \geq 0}$ is called the k -particle exclusion process on G : see Figure 1. In this paper we are interested in a natural extension of the exclusion process to hypergraphs.




Figure 1: Example transition of 3-particle exclusion process on K_5 . When the edge indicated rings, the single particle currently on that edge moves to the vertex at the other end of the edge.

Let $G = (V, E)$ be a finite connected hypergraph, where $E \subseteq \mathcal{P}(V)$, the power set of V . For each $e \in E$, denote by \mathcal{S}_e the symmetric group on the elements in e , and let $f_e : \mathcal{S}_e \rightarrow [0, 1]$ be a probability measure on \mathcal{S}_e . We write f to denote $\{f_e : e \in E\}$, the set of these measures. Consider k indistinguishable particles moving on V using rules 1. and 2. above and in addition:

- 3'. when an edge e rings, a permutation $\sigma \in \mathcal{S}_e$ is chosen according to f_e and every particle on a vertex in e moves simultaneously according to σ , i.e. a particle at vertex v moves to vertex $\sigma(v)$. (Note that as σ is a permutation, rule 1. is preserved.)

Setting $\eta_t(v) = 1$ if v is occupied at time t and 0 otherwise, we obtain a process $(\eta_t)_{t \geq 0}$ referred to as the k -particle exclusion process on (f, G) , or simply $\text{EX}(k, f, G)$: see Figure 2. (Note that if each edge $e \in E$ contains exactly two vertices, and f_e puts all of its mass on the transposition belonging to \mathcal{S}_e , then $\text{EX}(k, f, G)$ is just the k -particle exclusion process on the graph G , as above.)




Figure 2: Example transition of 3-particle exclusion process on a hypergraph with 5 vertices and 3 edges (indicated by the different shaded regions, i.e. here there are two edges of size 3 and one of size 4). When the edge containing four vertices rings, the two particles currently belonging to that edge are permuted.

Our main aim in this paper is to study the total-variation mixing time of $\text{EX}(k, f, G)$, and to establish an upper bound in terms of the mixing time of $\text{EX}(2, f, G)$. Recall that for a continuous-time Markov process X on a finite set Ω with transition probabilities $\{q_t(x, y)\}$ and equilibrium distribution π , the total variation ε -mixing time is defined as

$$T_X(\varepsilon) := \inf \left\{ t \geq 0 : \max_{x \in \Omega} \|q_t(x, \cdot) - \pi\|_{\text{TV}} \leq \varepsilon \right\}, \tag{1.1}$$

where $\|\cdot\|_{\text{TV}}$ is the total-variation norm.

In several parts of the proof it will be useful to consider the associated process where the k particles are distinguishable. Suppose the particles are labelled $1, \dots, k$ and set $\hat{\eta}_t(v)$ to be the label of the particle at vertex v at time t . If there is no particle at v at time t , set $\hat{\eta}_t(v) = 0$. The process $(\hat{\eta}_t)_{t \geq 0}$ is the k -particle interchange process on (f, G) , or simply $\text{IP}(k, f, G)$. Note that the exclusion process may be recovered from the interchange process simply by ‘forgetting’ the labels of the particles.

Throughout we will make the following assumptions about the hypergraph G and the set of measures f (with notation appearing below being formally defined in Section 2.1).

Assumption 1.1.

1. The hypergraph G is regular (every vertex has the same degree).
2. For every e , f_e is constant on the conjugacy classes of \mathcal{S}_e (i.e. in group-theoretic terms, f_e is a class function). That is, if σ_1 and σ_2 are elements from \mathcal{S}_e with the same cycle structure, then $f_e(\sigma_1) = f_e(\sigma_2)$.
3. For every e and each $v \in e$, $\sum_{\sigma \in \mathcal{S}_e} f_e(\sigma) \mathbf{1}_{\{\sigma(v)=v\}} \leq 1/5$. In other words, the probability (under f_e) of a vertex $v \in e$ being a fixed point of σ is at most $1/5$.
4. The interchange process $\text{IP}(k, f, G)$ is irreducible for any number of particles $k \in \{1, \dots, |V| - 1\}$.

These assumptions are more than enough to imply that the exclusion process is reversible and ergodic, with uniform stationary distribution. Although we state it as an assumption on f , the fourth assumption also implies that the underlying hypergraph G is connected. Our main theorem is the following:

Theorem 1.2. *There exists a universal constant $C > 0$ such that for every (f, G) satisfying Assumption 1.1 and every $k \in \{1, \dots, |V| - 1\}$ and $\varepsilon > 0$,*

$$T_{\text{EX}(k, f, G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\text{EX}(2, f, G)}(1/4).$$

Remark 1.3. In all further statements we implicitly assume that Assumption 1.1 holds.

Remark 1.4. The exclusion process on a hypergraph G with the edge set E consisting only of edges of size 2 or 3 exhibits the *negative correlation* property (which we shall discuss further in the sequel). As a result, for this subset of hypergraphs we can actually extend the main theorem, replacing the $T_{\text{EX}(2, f, G)}(1/4)$ appearing in the right-hand side by $T_{\text{EX}(1, f, G)}(1/4)$ (note that we later refer to $\text{EX}(1, f, G)$ as $\text{RW}(1, f, G)$, in recognition that the exclusion process with just one particle is equivalent to a single particle performing a random walk).

Remark 1.5. A simple example suffices to show that Theorem 1.2 is optimal in the sense that we cannot replace $T_{\text{EX}(2, f, G)}(1/4)$ on the right-hand side with $T_{\text{EX}(1, f, G)}(1/4)$ (even under our standing assumptions). Let $G = (V, E)$ with $V = \{1, 1', 2, 2', \dots, m, m'\}$ and $E = \{\{i, i', j, j'\} : i \neq j\}$. Suppose that $f_{\{1, 1', 2, 2'\}}(\sigma) = 1/6$ if σ is a cycle of size 4 (and otherwise $f_{\{1, 1', 2, 2'\}}(\sigma) = 0$). For $\{i, i', j, j'\} \neq \{1, 1', 2, 2'\}$, we set $f_{\{i, i', j, j'\}}(\sigma) = 1/3$ if σ is a composition of two disjoint transpositions (and otherwise $f_{\{i, i', j, j'\}}(\sigma) = 0$). It can be readily checked that this hypergraph satisfies Assumption 1.1. Note that there are $\binom{m}{2}$ edges, each ringing at rate 1. It is easy to see that a random walker mixes in time of order $1/m$ since each vertex is in order m edges. Now consider a 2-particle exclusion process started from $\{3, 3'\}$. Notice that up until the first time that both particles occupy vertices belonging to the edge $\{1, 1', 2, 2'\}$, if vertex i is occupied then vertex i' is also occupied. So the process cannot mix until the edge $\{1, 1', 2, 2'\}$ is visited by the particles. Regardless of where the particles are before this time, there are always two edges that

can ring which would bring the particles to the set $\{1, 1', 2, 2'\}$, and so this happens at rate 2; we conclude that it takes a time of order 1 for the 2-particle exclusion process to mix.

1.1 Motivation and connections with the literature

Our results contribute to the general question of when properties of a multi-particle system can be deduced from properties of a system with only a few particles. Arguably the most significant recent result in this area has come from Caputo, Liggett and Richthammer ([3]) who showed that the spectral gap of the interchange process on a graph is equal to the spectral gap of a random walker on the same graph, proving a conjecture of Aldous that had been open for 20 years. Proving results in this area is particularly important in applications since the large reduction in the size of the state space often makes it much easier to compute or estimate statistics.

While interacting particle system models (e.g. exclusion process, interchange process, voter model, contact process, zero range process) on graphs have received considerable attention, there has so far been little study of such processes on hypergraphs. Studying these processes on hypergraphs is very natural though, as hypergraphs allow simultaneous interactions of multiple particles, rather than only pair-wise interactions. One model for which its analogue on hypergraphs has been recently studied is the voter model ([4, 8]), for which various properties are considered, including the mixing time.

Any interchange process (with $k = |V|$) on a graph can be viewed as a card shuffle by transpositions, and there is now an extensive literature concerning mixing times of such shuffles. Notable examples include the top-to-random transposition shuffle (star graph; [6]), random-to-random transposition shuffle (complete graph; [5]) and nearest-neighbour transposition shuffle (the cycle; [9]). Of course, transposition shuffles are just one class of shuffle, and there is significant interest in mixing times of more general shuffles in which multiple cards are moved simultaneously. A large class of time-homogeneous shuffles can be represented as interchange processes on hypergraphs; recent examples can be found in [2] and [7].

Achieving tight bounds on the mixing time of an interacting particle system typically involves finding an argument tailored specifically to the model in question. If we care less about the specific constant multiple (at which mixing occurs) and instead focus on the order, a result of Oliveira ([15]) can prove particularly useful as a general way of bounding mixing times of exclusion processes:

Theorem 1.6 ([15]). *There exists a constant $C > 0$ such that for every connected weighted graph G and every $k \in \{1, \dots, |V| - 1\}$ and $\varepsilon \in (0, 1/2)$,*

$$T_{\text{EX}(k,G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\text{RW}(G)}(1/4),$$

where $T_{\text{RW}(G)}(1/4)$ is the mixing time of the random walk on G .

Our main result extends Theorem 1.6 to a class of hypergraphs. Furthermore, our results hold for a large class of measures acting on the symmetric group $\mathcal{S}_{|V|}$ which goes beyond the standard framework studied by previous authors, in which a conjugacy class is fixed and then sampled from uniformly (e.g. [13, 2]). Indeed, our measures f_e can vary dramatically between edges $e \in E$, and furthermore we do not require each f_e to be supported on a fixed conjugacy class.

1.2 Heuristics and structure of the proof

The proof of Theorem 1.2 depends on the size of the vertex set V . If $|V|$ is sufficiently small, the proof is fairly simple and we state the result as the following lemma:

Lemma 1.7. *There exists a constant $C > 0$ such that for every hypergraph $G = (V, E)$*

with $|V| < 36$, every f and every $k \in \{1, \dots, |V|/2\}$ and $\varepsilon > 0$,

$$T_{\text{EX}(k,f,G)}(\varepsilon) \leq C \log(1/\varepsilon) T_{\text{EX}(2,f,G)}(1/4).$$

On the other hand, the argument for $|V| \geq 36$ is much more intricate and is split into two parts, the first being the following lemma which is of independent interest (and is stronger than needed for our main theorem, as it relates to the interchange process):

Lemma 1.8. *There exists a constant $C > 0$ such that for every hypergraph $G = (V, E)$ with $|V| \geq 36$, every f and every $k \in \{1, \dots, |V|/2\}$ and $\varepsilon > 0$,*

$$T_{\text{IP}(k,f,G)}(\varepsilon) \leq C \log(|V|/\varepsilon) T_{\text{EX}(4,f,G)}(1/4).$$

Oliveira ([15]) proves his main result (bounding the mixing time of the k -particle exclusion process by the mixing time of a random walker) by first relating the mixing time of a k -particle interchange process to that of a 2-particle interchange process. Roughly speaking, this is possible due to the fact that any time an edge of the graph under consideration rings, at most two particles move under interchange, and so it is pairwise interactions that determine the mixing rate. This contrasts with the exclusion process on hypergraphs considered here, in which many particles can move at the same time. Nevertheless, a suitable adaptation of the techniques appearing in [15] provides the proof of Lemma 1.8.

Remark 1.9. Lemma 1.8 only holds when $|V|$ is sufficiently large and $k \leq |V|/2$. We cannot hope to remove these conditions and replace $\text{EX}(4, f, G)$ with $\text{EX}(2, f, G)$ in this statement, even for hypergraphs satisfying Assumption 1.1, as the following example illustrates. Let $G = (V, E)$ with $V = \{1, 2, 3\}$ and $E = \{V\}$, i.e. there is just a single edge which contains all three vertices in the hypergraph. Suppose that f gives probability $1 - \delta$ to the conjugacy class of 3-cycles, and probability δ to the class of transpositions. For δ sufficiently small this satisfies Assumption 1.1. The 2-particle interchange process cannot mix until a transposition is chosen (as half of the states cannot be reached before this time), whereas this event is not necessary for the 2-particle exclusion process to mix, and hence it is straightforward to see that as $\delta \rightarrow 0$ we have $T_{\text{IP}(2,f,G)}(1/4)/T_{\text{EX}(2,f,G)}(1/4) \rightarrow \infty$.

The second part of the proof for $|V| \geq 36$ requires showing that $T_{\text{EX}(4,f,G)}$ and $T_{\text{EX}(2,f,G)}$ are of the same order:

Lemma 1.10. *There exists a constant $\lambda > 0$ such that for any hypergraph G with $|V| \geq 36$, and f ,*

$$T_{\text{EX}(4,f,G)}(1/4) \leq \lambda T_{\text{EX}(2,f,G)}(1/4).$$

We now demonstrate that Theorem 1.2 follows simply from Lemmas 1.7, 1.8 and 1.10.

Proof of Theorem 1.2. The contraction principle (see [1]) gives

$$T_{\text{EX}(k,f,G)}(\varepsilon) \leq T_{\text{IP}(k,f,G)}(\varepsilon),$$

and so provided $k \leq |V|/2$, we have the result for $|V| \geq 36$ by Lemmas 1.8 and 1.10 and for $|V| < 36$ by Lemma 1.7. However, note that switching the roles of occupied and unoccupied vertices in $\text{EX}(k, f, G)$ yields the process $\text{EX}(|V| - k, f, G)$. It follows that

$$T_{\text{EX}(k,f,G)}(\varepsilon) = T_{\text{EX}(|V|-k,f,G)}(\varepsilon),$$

and so the proof of Theorem 1.2 is complete. □

We finish this section with a brief overview of the rest of the paper. In Section 2 we define formally the processes considered in this paper and present some preliminary results. In addition, we demonstrate that the negative correlation property, which is fundamental to the result in [15], fails to hold for the hypergraph setting. In Section 3 we prove Lemma 1.8 subject to the existence of a process with certain key properties that relate it to an interchange process: see Lemma 3.1 for the precise statement. This process is constructed in Section 6 and we show it has the desired properties in Section 7. Proving Lemma 3.1 is the most challenging (and technical) part of this paper.

In Section 4 we prove Lemma 1.10 by first characterizing every hypergraph as one of two types depending on how long it takes any two of four independent particles to meet.

We use some of the ideas developed in Section 4 to prove Lemma 1.7 in Section 5. A few of the more technical proofs required are included in two appendices.

2 Preliminaries

2.1 Random walks, exclusion and interchange processes

We formally define the main processes studied in this paper, $RW(f, G)$, $RW(k, f, G)$, $EX(k, f, G)$ and $IP(k, f, G)$, by explicitly stating their generators. In the next section we shall present a *graphical construction* of these processes, similar to that of Liggett ([12]) for the standard interchange and exclusion processes. This graphical construction will allow us to simultaneously define the processes on the same probability space, and thus directly compare them.

Recall \mathcal{S}_e as the group of permutations of elements in e . Our processes of interest evolve by the action of permutations from these groups. However, it will often be convenient to consider permutations as acting on V and we can easily do this by extending a permutation $\sigma_e \in \mathcal{S}_e$ to a permutation in \mathcal{S}_V by setting $\sigma_e(v) = v$ for all $v \notin e$. We can also consider such permutations as acting on a subset of V or on vectors with elements being distinct members of V . To do this we can define, for a set $A \subseteq V$, $\sigma_e(A) := \{\sigma_e(a) : a \in A\}$, and for a vector \mathbf{x} of k distinct elements of V we define $\sigma_e(\mathbf{x}) := (\sigma_e(\mathbf{x}(i)))_{i=1}^k$.

Set notation: For $k \in \mathbb{N}$ we define

$$\binom{V}{k} := \{A \subseteq V : |A| = k\},$$

and for a set $A \subseteq V$ we write

$$(A)_k := \{\mathbf{a} = (\mathbf{a}(1), \dots, \mathbf{a}(k)) \in A^k : \mathbf{a}(i) \neq \mathbf{a}(j) \forall i \neq j\}.$$

Generators: We now explicitly state the generators of the processes. For a hypergraph G and a suitable set of functions f , the simple random walk on G , $RW(f, G)$, is the continuous-time Markov chain with state space V and generator

$$U^{RW}h(u) = \sum_{e \in E} \sum_{\sigma \in \mathcal{S}_e} f_e(\sigma)(h(\sigma(u)) - h(u))$$

for all $u \in V$ and $h : V \rightarrow \mathbb{R}$.

We denote by $RW(k, f, G)$ the product of k independent random walkers on G . This process is the continuous-time Markov chain with state space V^k and generator

$$U^{RW(k)}h(\mathbf{u}) = \sum_{e \in E} \sum_{i=1}^k \sum_{\sigma \in \mathcal{S}_e} f_e(\sigma)(h(\mathbf{u}_{\sigma(\mathbf{u}(i))}^i) - h(\mathbf{u})),$$

for all $\mathbf{u} \in V^k$ and $h : V^k \rightarrow \mathbb{R}$, where

$$\mathbf{u}_v^i(j) = \begin{cases} \mathbf{u}(j) & j \neq i, \\ v & j = i. \end{cases}$$

The k -particle exclusion process $\text{EX}(k, f, G)$, is the continuous-time Markov chain with state space $\binom{V}{k}$ and generator

$$U^{\text{EX}}h(A) = \sum_{e \in E} \sum_{\sigma \in \mathcal{S}_e} f_e(\sigma)(h(\sigma(A)) - h(A)),$$

for all $A \in \binom{V}{k}$ and $h : \binom{V}{k} \rightarrow \mathbb{R}$.

The k -particle interchange process $\text{IP}(k, f, G)$, is the continuous-time Markov chain with state space $(V)_k$ and generator

$$U^{\text{IP}}h(\mathbf{x}) = \sum_{e \in E} \sum_{\sigma \in \mathcal{S}_e} f_e(\sigma)(h(\sigma(\mathbf{x})) - h(\mathbf{x})),$$

for all $\mathbf{x} \in (V)_k$ and $h : (V)_k \rightarrow \mathbb{R}$.

2.2 Graphical construction

We first construct an independent sequence of E -valued random variables $\{e_n\}_{n \in \mathbb{N}}$ such that each e_n is identically distributed with $\mathbb{P}[e_n = e] = 1/|E|$. Given the sequence $\{e_n\}_{n \in \mathbb{N}}$, let $\{\sigma_n\}_{n \in \mathbb{N}}$ be a sequence of permutations with $\sigma_n \in \mathcal{S}_{e_n}$ independently chosen and satisfying for each $n \in \mathbb{N}$, and $e \in E$, $\mathbb{P}[\sigma_n = \sigma | e_n = e] = f_e(\sigma)$. Now that we have the sequence of edges that ring and the permutations to apply, it remains to determine the update times of the processes.

Let Λ be a Poisson process of rate $|E|$ and for $0 < s < t$ denote by $\Lambda[s, t]$ the number of points of Λ in $[s, t]$. For every $0 < s < t$, we define a random permutation $I_{[s,t]} : V \rightarrow V$ associated with the time interval $[s, t]$ to be the composition of the permutations performed during this time; that is,

$$I_{[s,t]} = \sigma_{e_{\Lambda[0,t]}} \circ \sigma_{e_{\Lambda[0,t]-1}} \circ \dots \circ \sigma_{e_{\Lambda[0,s]+1}}.$$

We set $I_t := I_{[0,t]}$ for each $t > 0$, and $I_{(t,t]}$ to be the identity. Note (cf Proposition 3.2 of [15]) that

$$\mathcal{L}[I_{(s,t]}] = \mathcal{L}[I_{(s,t]}^{-1}], \tag{2.1}$$

where we write \mathcal{L} for the law of a process.

We can lift the functions $I_{[s,t]}$ to functions on $\binom{V}{k}$ and $(V)_k$ in the following way: for $A \in \binom{V}{k}$,

$$I_{[s,t]}(A) = \{I_{[s,t]}(a) : a \in A\},$$

and for $\mathbf{x} \in (V)_k$,

$$I_{[s,t]}(\mathbf{x}) = (I_{[s,t]}(\mathbf{x}(1)), \dots, I_{[s,t]}(\mathbf{x}(k))).$$

The following proposition is fundamental: its proof follows by inspection.

Proposition 2.1. *Fix $s > 0$. Then:*

1. *For each $u \in V$, the process $\{I_{[s,s+t]}(u)\}_{t \geq 0}$ is a realisation of $\text{RW}(f, G)$ initialised at u at time s . We shall often write this process simply as $(u_t^{\text{RW}})_{t \geq s}$.*
2. *For each $A \in \binom{V}{k}$, the process $\{I_{[s,s+t]}(A)\}_{t \geq 0}$ is a realisation of $\text{EX}(k, f, G)$ initialised at A at time s . We shall often write this process simply as $(A_t^{\text{EX}})_{t \geq s}$.*
3. *For each $\mathbf{x} \in (V)_k$, the process $\{I_{[s,s+t]}(\mathbf{x})\}_{t \geq 0}$ is a realisation of $\text{IP}(k, f, G)$ initialised at \mathbf{x} at time s . We shall often write this process simply as $(\mathbf{x}_t^{\text{IP}})_{t \geq s}$.*

2.3 Total variation and mixing times

There are several equivalent definitions of total variation that we shall make use of in this paper. Suppose μ and ν are two probability measures on the same finite set Ω . Then the *total variation distance* between these measures is defined as

$$\|\mu - \nu\|_{\text{TV}} := \max_{A \subset \Omega} (\mu(A) - \nu(A)) \tag{2.2}$$

$$= \sup_{f: \Omega \rightarrow [0,1]} \int f d\mu - \int f d\nu. \tag{2.3}$$

We shall also make extensive use of the following equivalent definition, which relates the total variation distance to couplings of μ and ν :

$$\|\mu - \nu\|_{\text{TV}} = \inf_{(X,Y)} \mathbb{P}[X \neq Y], \tag{2.4}$$

where the infimum is over all couplings (X, Y) of random variables with $X \sim \mu$ and $Y \sim \nu$. We recall a simple result bounding the total variation of product chains (see e.g. pg 59 of [10]): for $n \in \mathbb{N}$ and $1 \leq i \leq n$, let μ_i and ν_i be measures on a finite space Ω_i and define measures μ and ν on $\prod_{i=1}^n \Omega_i$ by $\mu := \prod_{i=1}^n \mu_i$ and $\nu := \prod_{i=1}^n \nu_i$. Then

$$\|\mu - \nu\|_{\text{TV}} \leq \sum_{i=1}^n \|\mu_i - \nu_i\|_{\text{TV}}. \tag{2.5}$$

Recall equation (1.1) as the definition of the mixing time of a continuous-time Markov process. We will require several general mixing-time bounds throughout this work, which we present here.

Proposition 2.2 ([10]). *Let X be a Markov process on a finite state space. Then for every $\varepsilon_1, \varepsilon_2 \in (0, 1/2)$,*

$$T_X(\varepsilon_2) \leq \left\lceil \frac{\log \varepsilon_2}{\log(2\varepsilon_1)} \right\rceil T_X(\varepsilon_1).$$

Proposition 2.3. *For any $m, n \in \mathbb{N}$,*

$$T_{\text{RW}(2^m, f, G)}(2^{-n}) \leq (n + m) T_{\text{RW}(f, G)}(1/4).$$

Proof. This follows by combining Proposition 2.2 with (2.5). □

Proposition 2.4 ([1]). *Let X be a Markov process on a finite state space Ω with symmetric transition rates. Then the equilibrium distribution is uniform over Ω and for all $0 < \varepsilon < 1/2$ and $t \geq 2T_X(\varepsilon)$,*

$$\mathbb{P}[X_t = \omega_2 | X_0 = \omega_1] \geq \frac{(1 - 2\varepsilon)^2}{|\Omega|},$$

for all $\omega_1, \omega_2 \in \Omega$.

2.4 Failure of negative correlation

We conclude this preliminary section with a quick example to demonstrate that the exclusion process on a hypergraph does not enjoy the negative correlation property satisfied by the exclusion process on a graph. We first recall the version of the negative correlation property of the exclusion process on a graph to which we refer, and whose proof may be found in [11]. Let $B \subset V$ and let $(A_t^{\text{EX}})_{t \geq 0}$ be a 2-particle exclusion process on a graph $G = (V, E)$ with $A = \{u, v\}$. Suppose $(u_t^{\text{RW}})_{t \geq 0}$ and $(v_t^{\text{RW}})_{t \geq 0}$ are two

independent realisations of $\text{RW}(1, f, G)$, started from u and v respectively. Then for every $t \geq 0$,

$$\mathbb{P} [A_t^{\text{EX}} \subseteq B] \leq \mathbb{P} [u_t^{\text{RW}} \in B] \mathbb{P} [v_t^{\text{RW}} \in B].$$

Now suppose $G = (V, E)$ is the hypergraph with $V = \{1, 2, 3, 4\}$ and $E = \{V\}$ (i.e. there is only one edge), and that f is concentrated uniformly on the six possible 4-cycles. Let $(A_t^{\text{EX}})_{t \geq 0}$ be a realisation of $\text{EX}(2, f, G)$, with $A = \{u, v\} = \{1, 2\}$, and let $B = \{3, 4\}$. We claim that there exist values of t such that

$$\mathbb{P} [u_t^{\text{RW}} \in B, v_t^{\text{RW}} \in B] < \mathbb{P} [A_t^{\text{EX}} = B]. \tag{2.6}$$

Indeed, since the event $\{u_t^{\text{RW}} \in B\}$ is less likely than seeing at least one incident in a unit-rate Poisson process by time t , we have

$$\mathbb{P} [u_t^{\text{RW}} \in B, v_t^{\text{RW}} \in B] = \mathbb{P} [u_t^{\text{RW}} \in B] \mathbb{P} [v_t^{\text{RW}} \in B] \leq (1 - e^{-t})^2.$$

On the other hand, the event $\{A_t^{\text{EX}} = B\}$ is at least as likely as the edge ringing exactly once by time t , with the chosen permutation satisfying $\sigma(\{1, 2\}) = \{3, 4\}$. That is,

$$\mathbb{P} [A_t^{\text{EX}} = B] \geq \frac{1}{3} t e^{-t}.$$

Inequality (2.6) is therefore satisfied for any $t < 0.33$.

3 From k -particle interchange to 4-particle exclusion

In this section we shall prove Lemma 1.8. Given a hypergraph with vertex set V and a $(k - 1)$ -tuple $\mathbf{z} \in (V)_{k-1}$, let

$$\mathbf{O}(\mathbf{z}) := \{\mathbf{z}(1), \dots, \mathbf{z}(k - 1)\}$$

be the (unordered) set of coordinates of \mathbf{z} and define a space

$$\Omega_k(V) := \{(\mathbf{z}, R, P, W) : \mathbf{z} \in (V)_{k-1}, \text{ and sets } \mathbf{O}(\mathbf{z}), R, P, W \text{ partition } V\}.$$

As we shall see, most of the work required to prove Lemma 1.8 is to show the existence of a certain Markov process having some key properties, which we outline in the Lemma 3.1 below. As we shall see in the sequel, this Markov process is very similar to the chameleon process used in [15] and it provides a way of tracking how mixed the k th particle is in a k -particle interchange process. The k th particle is replaced by three sets of coloured particles, R_t (red particles), P_t (pink particles) and W_t (white particles), with the colours informing the conditional distribution of the k th particle in the interchange process. A process $(\text{ink}_t^{\mathbf{x}}(b))_{t \geq 0}$ is defined for each vertex $b \in V$, which records the amount of redness at vertex b (equal to 1 if a red particle is at vertex b and 1/2 if a pink particle is at vertex b). We shall also define an event $\text{Fill}^{\mathbf{x}}$ as the event that all vertices unoccupied by the first $k - 1$ particles in the interchange process are eventually each occupied by a red particle in the chameleon process.

Lemma 3.1. *There exist constants c_1, c_2 and κ_1 such that for every regular hypergraph $G = (V, E)$ with $|V| \geq 36$, every f , every $k \in \{2, \dots, |V|/2\}$, every $\mathbf{x} = (\mathbf{z}, x) \in (V)_k$, and every realisation $(\mathbf{x}_t^{\text{IP}})_{t \geq 0}$ of $\text{IP}(k, f, G)$ started from state \mathbf{x} , there exists a continuous-time Markov process $(M_t)_{t \geq 0} := (\mathbf{z}_t^C, R_t, P_t, W_t)_{t \geq 0}$ with state-space $\Omega_k(V)$ defined on the same probability space as $(\mathbf{x}_t^{\text{IP}})_{t \geq 0}$ satisfying:*

1. $(\mathbf{z}_t^{\text{IP}})_{t \geq 0} = (\mathbf{z}_t^C)_{t \geq 0}$ almost surely;

2. for every $t \geq 0$ and $\mathbf{b} = (\mathbf{c}, b) \in (V)_k$,

$$\mathbb{P} [\mathbf{x}_t^{\text{IP}} = \mathbf{b}] = \mathbb{E} \left[\text{ink}_t^{\mathbf{x}}(b) \mathbf{1}_{\{\mathbf{z}_t^C = \mathbf{c}\}} \right],$$

where $\text{ink}_t^{\mathbf{x}}(b) := \mathbf{1}_{\{b \in R_t\}} + \frac{1}{2} \mathbf{1}_{\{b \in P_t\}}$;

3. for every $t \geq 0$ and $j \in \mathbb{N}$,

$$\mathbb{E} \left[1 - \frac{\text{ink}_t^{\mathbf{x}}}{|V| - k + 1} \mid \text{Fill}^{\mathbf{x}} \right] \leq c_1 \sqrt{|V|} e^{-c_2 j} + \exp \left\{ j - \frac{t}{\kappa_1 T_{\text{EX}}(4, f, G)(1/4)} \right\}$$

where $\text{ink}_t^{\mathbf{x}} := \sum_{b \in V} \text{ink}_t^{\mathbf{x}}(b)$ and $\text{Fill}^{\mathbf{x}} := \{\lim_{t \rightarrow \infty} \text{ink}_t^{\mathbf{x}} = |V| - k + 1\}$;

4. for every $t \geq 0$ and $\mathbf{c} \in (V)_{k-1}$,

$$\mathbb{P} [\{\mathbf{z}_t^C = \mathbf{c}\} \cap \text{Fill}^{\mathbf{x}}] = \frac{\mathbb{P} [\mathbf{z}_t^C = \mathbf{c}]}{|V| - k + 1}.$$

The proof of Lemma 3.1 is deferred to Section 7 and is a proof by construction: in Section 6 we will explicitly define a process and then proceed to show that it has the desired properties. We can now relate the total-variation distance between two realisations of $\text{IP}(k, f, G)$ to a certain expectation involving the amount of ink in the chameleon process M in the statement of Lemma 3.1. The following result is similar to Lemma 6.1 of [15]: we include a sketch of the proof to highlight the importance of constructing in Section 6 a chameleon process satisfying part 2 of Lemma 3.1.

Lemma 3.2. For every $t \geq 0$,

$$\sup_{\mathbf{x}, \mathbf{y} \in (V)_k} \|\mathcal{L}[\mathbf{x}_t^{\text{IP}}] - \mathcal{L}[\mathbf{y}_t^{\text{IP}}]\|_{\text{TV}} \leq 2k \sup_{\mathbf{w} \in (V)_k} \mathbb{E} \left[1 - \frac{\text{ink}_t^{\mathbf{w}}}{|V| - k + 1} \mid \text{Fill}^{\mathbf{w}} \right]$$

Proof. Fix $\mathbf{x} = (\mathbf{z}, x) \in (V)_k$ with $\mathbf{z} \in (V)_{k-1}$, and denote by \mathbf{x}_t^{IP} an interchange process started from \mathbf{x} . Let \tilde{x} be uniform from $V \setminus \mathbf{O}(\mathbf{z})$ and denote by $\tilde{\mathbf{x}}_t^{\text{IP}}$ an interchange process started from $\tilde{\mathbf{x}} = (\mathbf{z}, \tilde{x})$. Then for any $\mathbf{b} = (\mathbf{c}, b) \in (V)_k$,

$$\mathbb{P} [\tilde{\mathbf{x}}_t^{\text{IP}} = \mathbf{b}] = \frac{\mathbb{P} [\mathbf{z}_t^{\text{IP}} = \mathbf{c}]}{|V| - k + 1} = \frac{\mathbb{P} [\mathbf{z}_t^C = \mathbf{c}]}{|V| - k + 1} = \mathbb{P} [\{\mathbf{z}_t^C = \mathbf{c}\} \cap \text{Fill}^{\mathbf{x}}],$$

where the second and third equalities follow from parts 1 and 4 of Lemma 3.1, respectively. On the other hand, part 2 of Lemma 3.1 gives

$$\mathbb{P} [\mathbf{x}_t^{\text{IP}} = \mathbf{b}] = \mathbb{E}[\text{ink}_t^{\mathbf{x}}(b) \mathbf{1}_{\{\mathbf{z}_t^C = \mathbf{c}\}}] \geq \mathbb{E}[\text{ink}_t^{\mathbf{x}}(b) \mathbf{1}_{\{\{\mathbf{z}_t^C = \mathbf{c}\} \cap \text{Fill}^{\mathbf{x}}\}}].$$

Subtracting we obtain

$$\mathbb{P} [\tilde{\mathbf{x}}_t^{\text{IP}} = \mathbf{b}] - \mathbb{P} [\mathbf{x}_t^{\text{IP}} = \mathbf{b}] \leq \mathbb{E}[(1 - \text{ink}_t^{\mathbf{x}}(b)) \mathbf{1}_{\{\{\mathbf{z}_t^C = \mathbf{c}\} \cap \text{Fill}^{\mathbf{x}}\}}].$$

Hence

$$\begin{aligned} \|\mathcal{L}[\mathbf{x}_t^{\text{IP}}] - \mathcal{L}[\tilde{\mathbf{x}}_t^{\text{IP}}]\|_{\text{TV}} &\leq \sum_{(\mathbf{c}, b) \in (V)_k} \mathbb{E}[(1 - \text{ink}_t^{\mathbf{x}}(b)) \mathbf{1}_{\{\{\mathbf{z}_t^C = \mathbf{c}\} \cap \text{Fill}^{\mathbf{x}}\}}] \\ &= \mathbb{E}[(|V| - k + 1 - \text{ink}_t^{\mathbf{x}}) \mathbf{1}_{\{\text{Fill}^{\mathbf{x}}\}}] \\ &= \mathbb{E} \left[1 - \frac{\text{ink}_t^{\mathbf{x}}}{|V| - k + 1} \mid \text{Fill}^{\mathbf{x}} \right]. \end{aligned}$$

The result now follows by repeated application of the triangle inequality, as in the proof of Lemma 6.1 of [15]. \square

Proof of Lemma 1.8. We combine part 3 of Lemma 3.1 with Lemma 3.2 to give for every $t \geq 0$ and $j \in \mathbb{N}$,

$$\sup_{\mathbf{x}, \mathbf{y} \in (V)_k} \|\mathcal{L}[\mathbf{x}_t^{\text{IP}}] - \mathcal{L}[\mathbf{y}_t^{\text{IP}}]\|_{\text{TV}} \leq 2k \left\{ c_1 \sqrt{|V|} e^{-c_2 j} + \exp \left\{ j - \frac{t}{\kappa_1 T_{\text{EX}}(4, f, G)} (1/4) \right\} \right\},$$

for some universal positive constants c_1, c_2 and κ_1 . We choose

$$j = \left\lfloor \frac{t}{(1 + c_2) \kappa_1 T_{\text{EX}}(4, f, G)} (1/4) \right\rfloor,$$

which gives the bound (using $k \leq |V|$),

$$\sup_{\mathbf{x}, \mathbf{y} \in (V)_k} \|\mathcal{L}[\mathbf{x}_t^{\text{IP}}] - \mathcal{L}[\mathbf{y}_t^{\text{IP}}]\|_{\text{TV}} \leq c_3 |V|^{3/2} \exp \left\{ \frac{-c_2 t}{(1 + c_2) \kappa_1 T_{\text{EX}}(4, f, G)} (1/4) \right\},$$

for some positive c_3 . Therefore there exists a universal constant C such that for any $\varepsilon \in (0, 1/2)$ and $t > CT_{\text{EX}}(4, f, G) (1/4) \log(|V|/\varepsilon)$,

$$\sup_{\mathbf{x}, \mathbf{y} \in (V)_k} \|\mathcal{L}[\mathbf{x}_t^{\text{IP}}] - \mathcal{L}[\mathbf{y}_t^{\text{IP}}]\|_{\text{TV}} \leq \varepsilon. \quad \square$$

4 From 4-particle exclusion to 2-particle exclusion

In this section we shall prove Lemma 1.10. We begin by characterizing every connected hypergraph in terms of how long it takes two independent random walkers on the hypergraph to arrive onto the same edge, which then rings for one of the walkers – a time we shall refer to as the *meeting time* of the two walkers (note that we do not require the two walkers to actually occupy the same vertex). It will be useful to consider such times, as we will be able to couple two independent walkers with a 2-particle interchange process, up until this meeting time (see Proposition 4.7 for this statement).

Formalising this, for $\mathbf{y} \in V^2$, let $(\mathbf{y}_t^{\text{RW}})_{t \geq 0}$ be a realisation of $\text{RW}(2, f, G)$ with $\mathbf{y}_0^{\text{RW}} = \mathbf{y}$. Denote by Λ^1 and Λ^2 the Poisson processes used to generate the edge-ringing times for the two particles, and let $\{e_n^1\}_{n \in \mathbb{N}}$ and $\{e_n^2\}_{n \in \mathbb{N}}$ be the two sequences of edge-choices (all as in Section 2.2). Define $M^{\text{RW}}(\mathbf{y})$ to be the first time $\mathbf{y}_t^{\text{RW}}(1)$ and $\mathbf{y}_t^{\text{RW}}(2)$ are in the same edge which then rings in one of the processes:

$$M^{\text{RW}}(\mathbf{y}) := \inf \{t > 0 : \exists e \in \{e_{\Lambda^1[0, t]}^1, e_{\Lambda^2[0, t]}^2\} \text{ with } \mathbf{y}_{t-}^{\text{RW}}(1), \mathbf{y}_{t-}^{\text{RW}}(2) \in e\}. \quad (4.1)$$

Definition 4.1. We say that a hypergraph G is easy if

$$\sup_{\mathbf{y} \in V^2} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 10^{10} T_{\text{EX}}(2, f, G) (1/4)] \leq 1/1000.$$

Remark 4.2. We note that this definition is similar to Definition 4.1 in [15], from where we borrow the dichotomy “easy/non-easy”. However, for the case of hypergraphs, this characterisation does not reflect the associated difficulty of dealing with each case! One difference in the case of hypergraphs is that at the meeting time we cannot guarantee that the two independent walkers occupy the same site, and this results in the analysis being more challenging.

4.1 From 4-particle exclusion to 2-particle exclusion: easy hypergraphs

We present a preliminary result which shows that we can couple two k -particle exclusion processes initially sharing $k - 1$ occupied vertices such that, with positive probability, at the meeting time of the k th particles the two processes will agree, given

the k th particles meet on an edge of size at least 5 and the permutation chosen at this meeting time does not fix the k th particles.

Let Λ be a Poisson process of rate $2|E|$ (i.e. twice the usual rate), with associated edge-choices $\{e_n\}_{n \in \mathbb{N}}$ and permutations $\{\sigma_n\}_{n \in \mathbb{N}}$ as in Section 2.2. In addition, let $\{\theta_n\}_{n \in \mathbb{N}}$ be an i.i.d. sequence of Bernoulli(1/2) random variables: these will be used to thin the events of Λ and ensure that all particles are moving at the correct rate. We write $\hat{\Lambda}$ for the thinned Poisson process obtained from Λ by removing all points corresponding to $\theta_n = 0$. Let \hat{I}_t be constructed from $\hat{\Lambda}$ as in Section 2.2. We make this modification as it allows us to more easily compare a certain time to the meeting time of two independent random walkers as defined in (4.1).

Lemma 4.3. *Let $D \in \binom{V}{k-1}$, $a, b \in V \setminus D$ and $A = D \cup \{a\}$, $B = D \cup \{b\}$. Let $(A_t^{\text{EX}})_{t \geq 0}$ and $(B_t^{\text{EX}})_{t \geq 0}$ be two realisations of $EX(k, f, G)$ started from A and B respectively and evolving according to \hat{I}_t . Let*

$$\tau_{a,b} := \inf\{t \geq 0 : \hat{I}_t(a), \hat{I}_t(b) \in e_{\Lambda[0,t]}\},$$

and write $e_{a,b}$ for $e_{\Lambda[0,\tau_{a,b}]}$ and $\sigma_{a,b}$ for $\sigma_{\Lambda[0,\tau_{a,b}]}$. Write also a^* for $\hat{I}_{[0,\tau_{a,b})}(a)$ and b^* for $\hat{I}_{[0,\tau_{a,b})}(b)$. Then there exist two other realisations of $EX(k, f, G)$ denoted $(A_t^{\text{EX}})_{t \geq 0}$ and $(B_t^{\text{EX}})_{t \geq 0}$ which start and evolve identically to $(A_t^{\text{EX}})_{t \geq 0}$ and $(B_t^{\text{EX}})_{t \geq 0}$ respectively up to time $\tau_{a,b}$ — but which satisfy, on the event

$$\{\sigma_{a,b}(a^*) \neq a^*\} \cap \{|e_{a,b}| > 4\},$$

with probability at least $\frac{2}{25}$, $A_t^{\text{EX}} = B_t^{\text{EX}}$ for all $t \geq \tau_{a,b}$.

Proof. We define two events which will be used to determine the coupling strategy of the processes $(A_t^{\text{EX}})_{t \geq 0}$ and $(B_t^{\text{EX}})_{t \geq 0}$ at time $\tau_{a,b}$:

$$\begin{aligned} J_1(\sigma_{a,b}) &= \{\sigma_{a,b}(a^*) \notin \{a^*, b^*\}, \sigma_{a,b}(b^*) \notin \{a^*, b^*\}\} \\ J_2(\sigma_{a,b}) &= J_1(\sigma_{a,b}) \cap \left\{ \left| \hat{I}_{\tau_{a,b}-}(D) \cap \sigma_{a,b}(a^*) \right| = \left| \hat{I}_{\tau_{a,b}-}(D) \cap \sigma_{a,b}(b^*) \right| \right\}. \end{aligned}$$

In words, $J_1(\sigma_{a,b})$ is the event that the permutation $\sigma_{a,b}$ moves the set of two ‘special’ particles (those initially at vertices a and b) to a new set of positions; event $J_2(\sigma_{a,b})$ further specifies that the two positions to which $\sigma_{a,b}$ moves the special particles should either both contain another particle (i.e. one of the already-matched $k - 1$ particles) or both be empty.

With this notation in place, we can describe the coupling at time $\tau_{a,b}$:

- (i) if $\theta_{a,b} = 0$ then we do not update the processes at time $\tau_{a,b}$;
- (ii) if $\theta_{a,b} = 1$ but event $J_2(\sigma_{a,b})$ fails to hold, then we apply permutation $\sigma_{a,b}$ in both processes;
- (iii) if $\theta_{a,b} = 1$ and event $J_2(\sigma_{a,b})$ holds, we update the ‘A’ process with permutation $\sigma_{a,b}$ and the ‘B’ process with permutation $\bar{\sigma}_{a,b}$, where

$$\bar{\sigma}_{a,b} = \sigma_{a,b} \circ (\sigma_{a,b}(a^*) \sigma_{a,b}(b^*)) (\sigma_{a,b}^2(a^*) \sigma_{a,b}^2(b^*)).$$

(Here and throughout we use the convention that composition of permutations corresponds to multiplication on the right: $\sigma \circ \rho = \rho\sigma$.)

Figure 3 demonstrates the relationship between $\sigma_{a,b}$ and $\bar{\sigma}_{a,b}$ in the simple case where $\sigma_{a,b}$ is a single cycle. To show that this is a valid coupling, it suffices to show that in case (iii) the permutation $\bar{\sigma}_{a,b}$ belongs to the same conjugacy class as $\sigma_{a,b}$, and that there is a




Figure 3: The left/right image shows the state of the process on edge $e_{a,b}$ at time $\tau_{a,b-}$ in the ‘A’/‘B’ process. Also indicated are the permutations $\sigma_{a,b}$ and $\bar{\sigma}_{a,b}$ which are to be applied in case (iii).

bijection between the two permutations. By inspection, the cyclic decomposition of $\bar{\sigma}_{a,b}$ is obtained from that of $\sigma_{a,b}$ just by exchanging the elements $\sigma_{a,b}(a^*)$ and $\sigma_{a,b}(b^*)$, and so both permutations belong to the same conjugacy class. Moreover, there is a bijection between them since

$$\sigma_{a,b}(a^*) = \bar{\sigma}_{a,b}(b^*) \quad \text{and} \quad \sigma_{a,b}(b^*) = \bar{\sigma}_{a,b}(a^*),$$

and so $J_1(\sigma_{a,b}) = J_1(\bar{\sigma}_{a,b})$ and $J_2(\sigma_{a,b}) = J_2(\bar{\sigma}_{a,b})$.

Furthermore, it follows from the above analysis that our coupling strategy in case (iii) gives $\sigma_{a,b}(a^*) = \bar{\sigma}_{a,b}(b^*)$, and furthermore, $\sigma_{a,b}(\hat{I}_{\tau_{a,b-}}(D)) = \bar{\sigma}_{a,b}(\hat{I}_{\tau_{a,b-}}(D))$. Thus in order to complete this proof, we need to show that

$$\mathbb{P}[\theta_{a,b} = 1, J_2(\sigma_{a,b}) \mid \sigma_{a,b}(a^*) \neq a^*, \{|e_{a,b}| > 4\}] \geq 2/25.$$

We have

$$\begin{aligned} & \mathbb{P}[J_1(\sigma_{a,b}) \mid \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4] \\ &= \mathbb{P}[\sigma_{a,b}(a^*) \notin \{a^*, b^*\} \mid \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4] \\ & \quad \cdot \mathbb{P}[\sigma_{a,b}(b^*) \notin \{a^*, b^*\} \mid \sigma_{a,b}(a^*) \notin \{a^*, b^*\}, |e_{a,b}| > 4]. \end{aligned}$$

Using parts 2 and 3 of Assumption 1.1) this becomes

$$\begin{aligned} & \mathbb{P}[J_1(\sigma_{a,b}) \mid \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4] \\ & \geq \frac{4}{5} \left(1 - \mathbb{P}[\sigma_{a,b}(b^*) = b^*] - \frac{1 - \mathbb{P}[\sigma_{a,b}(b^*) = b^*]}{4} \right) \geq \frac{12}{25}. \end{aligned}$$

Therefore,

$$\begin{aligned} & \mathbb{P}[\theta_{a,b} = 1, J_2(\sigma_{a,b}) \mid \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4] \\ &= \frac{1}{2} \mathbb{P}[J_1(\sigma_{a,b}) \mid \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4] \\ & \quad \cdot \mathbb{P}[J_2(\sigma_{a,b}) \mid J_1(\sigma_{a,b}), \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4] \\ & \geq \frac{6}{25} \mathbb{P}[J_2(\sigma_{a,b}) \mid J_1(\sigma_{a,b}), \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4]. \end{aligned}$$

But conditioned on $J_1(\sigma_{a,b})$ and $\{\sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4\}$ both holding, $J_2(\sigma_{a,b})$ is the event that two of the positions in $e_{a,b}$ not containing a^* or b^* either both contain a matched particle or are both empty; since $|e_{a,b}| \geq 5$ this probability is at least $1/3$, thanks to part 2 of Assumption 1.1, and so our proof is complete. \square

We now present the main result of this section.

Lemma 4.4. *There exists $\kappa > 0$ such that for any easy hypergraph G , any f and $0 < \varepsilon < 1/2$,*

$$T_{\text{EX}(k,f,G)}(\varepsilon) \leq \kappa \log(1/\varepsilon) T_{\text{EX}(k-1,f,G)}(1/4),$$

for any $3 \leq k \leq |V|/2$ if $|V| < 36$ and any $k \in \{3, 4\}$ if $|V| \geq 36$.

In this section we will make use of this lemma only for the case $|V| \geq 36$, but this result will later be used in its full form when dealing with the case of $|V| < 36$: see Section 5. The proof uses a coupling argument for two realisations of $\text{EX}(k, f, G)$.

Proof. For $U = \{u_1, \dots, u_k\}$, $W = \{w_1, \dots, w_k\} \in \binom{V}{k}$, let $(U_t^{\text{EX}})_{t \geq 0}$ and $(W_t^{\widetilde{\text{EX}}})_{t \geq 0}$ be two realisations of $\text{EX}(k, f, G)$ started from U and W respectively. We define the two processes on a common probability space, and will show how to couple them in such a way that we can lower-bound the probability that $U_{\kappa T}^{\text{EX}} = W_{\kappa T}^{\widetilde{\text{EX}}}$ for some $\kappa > 0$ to be determined, where $T := T_{\text{EX}(k-1,f,G)}(1/4)$. The result will then follow by applying (2.4).

We begin by allowing the two processes to evolve independently up to time $10T$. Then, for any $S \subset \binom{V}{k}$ and $t \geq 0$, we have

$$\begin{aligned} \mathbb{P}[U_{10T+t}^{\text{EX}} \in S] - \mathbb{P}[W_{10T+t}^{\widetilde{\text{EX}}} \in S] &= \mathbb{E} \left[\mathbb{P}[U_{10T+t}^{\text{EX}} \in S | U_{10T}^{\text{EX}}] - \mathbb{P}[W_{10T+t}^{\widetilde{\text{EX}}} \in S | W_{10T}^{\widetilde{\text{EX}}}] \right] \\ &\leq \mathbb{E} \left[\|\mathcal{L}[U_{10T+t}^{\text{EX}} | U_{10T}^{\text{EX}}] - \mathcal{L}[W_{10T+t}^{\widetilde{\text{EX}}} | W_{10T}^{\widetilde{\text{EX}}}]\|_{\text{TV}} \right], \end{aligned}$$

where the inequality follows from (2.2). Maximizing over S and again using (2.2) gives

$$\|\mathcal{L}[U_{10T+t}^{\text{EX}}] - \mathcal{L}[W_{10T+t}^{\widetilde{\text{EX}}}\|_{\text{TV}} \leq \mathbb{E} \left[\|\mathcal{L}[U_{10T+t}^{\text{EX}} | U_{10T}^{\text{EX}}] - \mathcal{L}[W_{10T+t}^{\widetilde{\text{EX}}} | W_{10T}^{\widetilde{\text{EX}}}\|_{\text{TV}} \right]. \quad (4.2)$$

By the Markov property, for any $A, B \in \binom{V}{k}$,

$$\begin{aligned} \|\mathcal{L}[U_{10T+t}^{\text{EX}} | U_{10T}^{\text{EX}} = A] - \mathcal{L}[W_{10T+t}^{\widetilde{\text{EX}}} | W_{10T}^{\widetilde{\text{EX}}} = B]\|_{\text{TV}} &= \|\mathcal{L}[A_t^{\text{EX}}] - \mathcal{L}[B_t^{\widetilde{\text{EX}}}\|_{\text{TV}} \\ &\leq \mathbb{P}[A_t^{\text{EX}} \neq B_t^{\widetilde{\text{EX}}}], \end{aligned} \quad (4.3)$$

for any coupling of $(A_t^{\text{EX}})_{t \geq 0}$ and $(B_t^{\widetilde{\text{EX}}})_{t \geq 0}$, by (2.4), and where $\mathcal{L}[\cdot]$ denotes a conditional law.

Recall from Section 2.2 the construction of the permutation I_t for each $t \geq 0$. For any $A, B \in \binom{V}{k}$, let a and b be two uniformly and independently chosen elements of A and B , respectively. Given a , consider now the k -particle process $(A_t^a)_{t \geq 0} = (I_t(a), I_t(A \setminus \{a\}))_{t \geq 0}$ which evolves in the same way as the exclusion process begun at A , but with the label of the particle started from position a being tracked. Thus $(A_t^a)_{t \geq 0}$ can be thought of as something ‘between’ an exclusion process (in which no labels are tracked) and an interchange process (in which all labels are tracked). It’s clear that the $k - 1$ particles initially at vertices in $A \setminus \{a\}$ behave marginally as an exclusion process, while the particle started from a behaves (again marginally) as a random walk on G . Furthermore, the exclusion process $(A_t^{\text{EX}})_{t \geq 0}$ can be recovered from $(A_t^a)_{t \geq 0}$ simply by ‘forgetting’ which position is occupied by the ‘special’ particle starting from a , i.e. $A_t^{\text{EX}} = \{I_t(a), I_t(A \setminus \{a\})\}$. In a similar manner, for given b and another permutation-valued process $(\tilde{I}_t)_{t \geq 0}$, we also define the process $(\tilde{B}_t^b)_{t \geq 0} = (\tilde{I}_t(b), \tilde{I}_t(B \setminus \{b\}))_{t \geq 0}$.

Over the time period $[0, 10T]$ we couple the processes $(A_t^a)_{t \geq 0}$ and $(\tilde{B}_t^b)_{t \geq 0}$ using a maximal coupling of the $(k - 1)$ -particle exclusion processes $I_t(A \setminus \{a\})$ and $\tilde{I}_t(B \setminus \{b\})$. (Recall that a maximal coupling is one which achieves equality in the coupling inequality (2.4). This maximal coupling is actually more than is needed here; we will only be interested in the state of the processes at time $10T$.) By Proposition 2.2 we have

$$T_{\text{EX}(k-1,f,G)}(1/500) \leq \left\lceil \frac{\log\left(\frac{1}{500}\right)}{\log\left(\frac{1}{2}\right)} \right\rceil T < 10T. \tag{4.4}$$

Given the choice of a and b , let $F_{a,b}$ denote the event that the other $(k - 1)$ particles have coupled by time $10T$, i.e. $F_{a,b} = \{I_{10T}(A \setminus \{a\}) = \tilde{I}_{10T}(B \setminus \{b\})\}$. Using this maximal coupling it follows from (4.4) that $\mathbb{P}[F_{a,b}] \geq 499/500$. Combining this with equations (4.2) and (4.3) we see that for any $K \in \mathbb{N}$,

$$\begin{aligned} & \|\mathcal{L}[U_{(20+K)T}^{\text{EX}}] - \mathcal{L}[W_{(20+K)T}^{\text{EX}}]\|_{\text{TV}} \\ & \leq \sum_{A,B \in \binom{V}{k}} \mathbb{P}\left[U_{10T}^{\text{EX}} = A, W_{10T}^{\text{EX}} = B\right] \mathbb{P}\left[A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\text{EX}}\right] \\ & = \sum_{A,B \in \binom{V}{k}} \mathbb{P}\left[U_{10T}^{\text{EX}} = A\right] \mathbb{P}\left[W_{10T}^{\text{EX}} = B\right] \mathbb{P}\left[A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\text{EX}}\right] \\ & \leq \sum_{A,B \in \binom{V}{k}} \mathbb{P}\left[U_{10T}^{\text{EX}} = A\right] \mathbb{P}\left[W_{10T}^{\text{EX}} = B\right] \\ & \quad \cdot \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \left(1 - \mathbb{P}[F_{a,b}] + \mathbb{P}\left[A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\text{EX}}, F_{a,b}\right]\right) \\ & \leq \frac{1}{500} + \sum_{A,B \in \binom{V}{k}} \mathbb{P}\left[U_{10T}^{\text{EX}} = A\right] \mathbb{P}\left[W_{10T}^{\text{EX}} = B\right] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P}\left[A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\text{EX}}, F_{a,b}\right], \end{aligned} \tag{4.5}$$

where the equality is thanks to the independence of U and W over $[0, 10T]$.

From (4.5) we see that we now need to upper bound the probability that $(A_t^{\text{EX}})_{t \geq 0}$ and $(B_t^{\text{EX}})_{t \geq 0}$ do not agree by time $(10 + K)T$, on the event $F_{a,b}$. As pointed out above, this event is equivalent (on $F_{a,b}$) to the *locations* of the k particles in $A_{(10+K)T}^a$ and $\tilde{B}_{(10+K)T}^b$ not agreeing.

We shall bound this probability by coupling the processes $(A_{10T+t}^a)_{t \geq 0}$ and $(\tilde{B}_{10T+t}^b)_{t \geq 0}$ in the following manner. Recall the Poisson process Λ of rate $2|E|$ at the start of Section 4.1 with associated edge-choices $\{e_n\}_{n \in \mathbb{N}}$, permutations $\{\sigma_n\}_{n \in \mathbb{N}}$, and Bernoulli $(1/2)$ random variables $\{\theta_n\}_{n \in \mathbb{N}}$ (used to thin the events of Λ). Prior to a time $\tau_{a,b}$ defined below we evolve $(A_{10T+t}^a)_{t \geq 0}$ and $(\tilde{B}_{10T+t}^b)_{t \geq 0}$ by applying permutation σ_n to edge e_n (in both processes) at the n^{th} incident time of Λ if and only if $\theta_n = 1$, so formally we have for each $0 \leq t < \tau_{a,b}$,

$$A_{10T+t}^a = \hat{I}_t(I_{10T}(a), I_{10T}(A \setminus \{a\})), \quad \tilde{B}_{10T+t}^b = \hat{I}_t(\tilde{I}_{10T}(b), \tilde{I}_{10T}(B \setminus \{b\})).$$

Note that, since we use a common set of innovations over the period $[10T, 10T + \tau_{a,b})$, on event $F_{a,b}$ we have $D := \hat{I}_{\tau_{a,b}-}(I_{10T}(A \setminus \{a\})) = \hat{I}_{\tau_{a,b}-}(\tilde{I}_{10T}(B \setminus \{b\}))$; that is, the locations of the $k - 1$ unlabelled particles of A^a and \tilde{B}^b still agree at time $\tau_{a,b}-$. By the Markov property, on event $F_{a,b}$ we can thus write

$$A_{10T+t}^a = \hat{I}_t(I_{10T}(a), D), \quad \tilde{B}_{10T+t}^b = \hat{I}_t(\tilde{I}_{10T}(b), D).$$

We define $\tau_{a,b}$ to be the first time that the ‘special’ particles initially at a and b are in a common edge which then rings (note this has a slightly different definition from $\tau_{a,b}$ defined in the statement of Lemma 4.3):

$$\tau_{a,b} := \inf\{t \geq 0 : \hat{I}_t(I_{10T}(a)), \hat{I}_t(\tilde{I}_{10T}(b)) \in e_{\Lambda[0,t]}\}.$$

Note that the processes $(\hat{I}_t(I_{10T}(a)))_{t \geq 0}$ and $(\hat{I}_t(\tilde{I}_{10T}(b)))_{t \geq 0}$ when viewed marginally behave as independent random walks over the period $[0, \tau_{a,b})$, and so $\tau_{a,b}$ has the same distribution as the meeting time $M^{\text{RW}}(I_{10T}(a), \tilde{I}_{10T}(b))$ in (4.1).

To determine how to couple the processes at time $\tau_{a,b}$ we partition the probability space according to the following four sets (for some $K \in \mathbb{N}$ which is yet to be determined), denoting $\hat{I}_{\tau_{a,b}-}(I_{10T}(a))$ by a^* for ease of readability:

$$\begin{aligned} E_{a,b}^1 &:= \{\tau_{a,b} > KT\}, \\ E_{a,b}^2 &:= \{\tau_{a,b} \leq KT, \sigma_{a,b}(a^*) = a^*\}, \\ E_{a,b}^3 &:= \{\tau_{a,b} \leq KT, \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4\}, \\ E_{a,b}^4 &:= \{\tau_{a,b} \leq KT, \sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| \leq 4\}. \end{aligned}$$

For the first two cases, we shall not specify the coupling, as it does not matter how we update the processes at time $\tau_{a,b}$. First, for the case of $E_{a,b}^1$, we have

$$\begin{aligned} &\sum_{A,B \in \binom{V}{k}} \mathbb{P}[U_{10T}^{\text{EX}} = A] \mathbb{P}[W_{10T}^{\text{EX}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P}[A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\text{EX}}, F_{a,b}, E_{a,b}^1] \\ &\leq \sum_{A,B \in \binom{V}{k}} \mathbb{P}[U_{10T}^{\text{EX}} = A] \mathbb{P}[W_{10T}^{\text{EX}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P}[E_{a,b}^1] \leq \max_{a,b} \mathbb{P}[E_{a,b}^1]. \end{aligned} \quad (4.6)$$

Second, for the case of $E_{a,b}^2$, we have

$$\begin{aligned} &\sum_{A,B \in \binom{V}{k}} \mathbb{P}[U_{10T}^{\text{EX}} = A] \mathbb{P}[W_{10T}^{\text{EX}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P}[A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\text{EX}}, F_{a,b}, E_{a,b}^2] \\ &\leq \sum_{A,B \in \binom{V}{k}} \mathbb{P}[U_{10T}^{\text{EX}} = A] \mathbb{P}[W_{10T}^{\text{EX}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P}[\sigma_{a,b}(a^*) = a^*] \\ &= \sum_{a,b \in V} \frac{1}{k^2} \mathbb{P}[\sigma_{a,b}(a^*) = a^*] \mathbb{P}[a \in U_{10T}^{\text{EX}}] \mathbb{P}[b \in W_{10T}^{\text{EX}}] \\ &\leq \frac{k^2}{|V|^2} \sum_{a,b \in V} \frac{1}{k^2} \mathbb{P}[\sigma_{a,b}(a^*) = a^*] + \|\mathcal{L}[(a,b)_{10T}^{\text{RW}}] - \text{Unif}(V^2)\|_{\text{TV}} \end{aligned} \quad (4.7)$$

$$\leq \frac{1}{|V|^2} \sum_{a,b \in V} \mathbb{P}[\sigma_{a,b}(a^*) = a^*] + \frac{2}{500}, \quad (4.8)$$

where the penultimate inequality uses (2.1) and (2.3) and the last inequality uses (2.5), (4.4) and the contraction principle.

Third, conditioned on the event $E_{a,b}^3$ and $F_{a,b}$, by Lemma 4.3 we can couple the

processes so that $\hat{I}_{\tau_{a,b}}(I_{10T}(A)) = \hat{I}_{\tau_{a,b}}(\tilde{I}_{10T}(B))$ with probability at least $2/25$, giving

$$\begin{aligned} & \sum_{A,B \in \binom{V}{k}} \mathbb{P} [U_{10T}^{\text{EX}} = A] \mathbb{P} [W_{10T}^{\widetilde{\text{EX}}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P} [A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\widetilde{\text{EX}}}, F_{a,b}, E_{a,b}^3] \\ & \leq \sum_{A,B \in \binom{V}{k}} \mathbb{P} [U_{10T}^{\text{EX}} = A] \mathbb{P} [W_{10T}^{\widetilde{\text{EX}}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \frac{23}{25} \mathbb{P} [\sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4] \\ & \leq \frac{2}{500} + \frac{23}{25|V|^2} \sum_{a,b \in V} \mathbb{P} [\sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| > 4], \end{aligned} \tag{4.9}$$

where the last inequality is obtained in the same way as (4.8).

Our fourth and final case to consider is $E_{a,b}^4$: on this event a simple case-by-case analysis (sketched in Appendix A) shows that as long as there are no other (already matched) particles on edge $e_{a,b}$ at time $\tau_{a,b}$ - (i.e. $|\hat{I}_{\tau_{a,b}}(I_{10T}(A)) \cap e_{a,b}| = 1$), there exists a bijection between permutations $\sigma_{a,b}$ and $\tilde{\sigma}_{a,b}$ such that $\tilde{\sigma}_{a,b}$ is a permutation with the same cycle structure as $\sigma_{a,b}$, and such that with probability at least $1/2$

$$\sigma_{a,b}(\hat{I}_{\tau_{a,b}}(I_{10T}(a))) = \tilde{\sigma}_{a,b}(\hat{I}_{\tau_{a,b}}(\tilde{I}_{10T}(b))).$$

That is, in this situation we are able to make the locations of all k particles of $A_{(10+K)T}^{\text{EX}}$ and $B_{(10+K)T}^{\widetilde{\text{EX}}}$ agree with probability at least $1/2$:

$$\mathbb{P} [A_{(10+K)T}^{\text{EX}} = B_{(10+K)T}^{\widetilde{\text{EX}}}, F_{a,b}, E_{a,b}^4] \geq \frac{1}{2} \mathbb{P} [|\hat{I}_{\tau_{a,b}}(I_{10T}(A)) \cap e_{a,b}| = 1, F_{a,b}, E_{a,b}^4].$$

We use a union bound to control the probability of the complement and write c^* for $\hat{I}_{\tau_{a,b}}(I_{10T}(c))$ and b^* for $\hat{I}_{\tau_{a,b}}(\tilde{I}_{10T}(b))$. We have

$$\begin{aligned} & \sum_{A,B \in \binom{V}{k}} \mathbb{P} [U_{10T}^{\text{EX}} = A] \mathbb{P} [W_{10T}^{\widetilde{\text{EX}}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P} [A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\widetilde{\text{EX}}}, F_{a,b}, E_{a,b}^4] \\ & \leq \sum_{A,B \in \binom{V}{k}} \mathbb{P} [U_{10T}^{\text{EX}} = A] \mathbb{P} [W_{10T}^{\widetilde{\text{EX}}} = B] \\ & \quad \cdot \sum_{\substack{a \in A \\ b \in B}} \frac{1}{2k^2} \left(\mathbb{P} [E_{a,b}^4] + \mathbb{P} [|\hat{I}_{\tau_{a,b}}(I_{10T}(A)) \cap e_{a,b}| > 1, a^* \neq b^*, E_{a,b}^4] \right) \\ & \leq \sum_{A,B \in \binom{V}{k}} \mathbb{P} [U_{10T}^{\text{EX}} = A] \mathbb{P} [W_{10T}^{\widetilde{\text{EX}}} = B] \\ & \quad \cdot \sum_{\substack{a \in A \\ b \in B}} \frac{1}{2k^2} \left(\mathbb{P} [E_{a,b}^4] + \sum_{c \in A \setminus \{a\}} \mathbb{P} [c^* \in e_{a,b}, c^* \neq b^*, a^* \neq b^*, E_{a,b}^4] \right) \\ & = \sum_{a,b \in V} \sum_{c \neq a} \frac{1}{2k^2} \left(\frac{\mathbb{P} [E_{a,b}^4]}{k-1} + \mathbb{P} [c^* \in e_{a,b}, c^* \neq b^*, a^* \neq b^*, E_{a,b}^4] \right) \\ & \quad \cdot \sum_{\substack{A \supset \{a,c\} \\ B \supset \{b\}}} \mathbb{P} [U_{10T}^{\text{EX}} = A] \mathbb{P} [W_{10T}^{\widetilde{\text{EX}}} = B]. \end{aligned} \tag{4.10}$$

We now upper bound this using (2.1) and (2.3) (using the same method as in (4.7)). This

gives the following upper bound for (4.10):

$$\begin{aligned} & \frac{2}{500} + \frac{k^2(k-1)}{|V|^2(|V|-1)} \sum_{a,b \in V} \sum_{c \neq a} \frac{1}{2k^2} \left(\frac{\mathbb{P}[E_{a,b}^4]}{k-1} + \mathbb{P}[c^* \in e_{a,b}, c^* \neq b^*, a^* \neq b^*, E_{a,b}^4] \right) \\ & \leq \frac{2}{500} + \frac{1}{2|V|^2} \sum_{a,b \in V} \mathbb{P}[E_{a,b}^4] \\ & \quad + \frac{k-1}{2|V|^2(|V|-1)} \sum_{a,b \in V} \sum_{c \neq a} \mathbb{P}[c^* \in e_{a,b}, c^* \neq b^*, a^* \neq b^*, E_{a,b}^4] \\ & \leq \frac{2}{500} + \frac{1}{2|V|^2} \sum_{a,b \in V} \mathbb{P}[E_{a,b}^4] + \frac{k-1}{|V|^2(|V|-1)} \sum_{a,b \in V} \mathbb{P}[E_{a,b}^4], \end{aligned}$$

since, on the event $E_{a,b}^4$, the size of edge $e_{a,b}$ is at most four and so on the event $\{a^* \neq b^*\}$ for any choice of $e_{a,b}$ there are only two possibilities for the value of c (since $c^* \notin \{a^*, b^*\} \subset e_{a,b}$). This gives

$$\begin{aligned} & \sum_{A,B \in \binom{V}{k}} \mathbb{P}[U_{10T}^{\text{EX}} = A] \mathbb{P}[W_{10T}^{\widetilde{\text{EX}}} = B] \sum_{\substack{a \in A \\ b \in B}} \frac{1}{k^2} \mathbb{P}[A_{(10+K)T}^{\text{EX}} \neq B_{(10+K)T}^{\widetilde{\text{EX}}}, F_{a,b}, E_{a,b}^4] \\ & \leq \frac{2}{500} + \frac{1}{|V|^2} \left(\frac{1}{2} + \frac{k-1}{|V|-1} \right) \sum_{a,b \in V} \mathbb{P}[\sigma_{a,b}(a^*) \neq a^*, |e_{a,b}| \leq 4]. \end{aligned} \tag{4.11}$$

We now combine the bounds in (4.5), (4.6), (4.8), (4.9) and (4.11) to see that

$$\begin{aligned} \|\mathcal{L}[U_{(20+K)T}^{\text{EX}}] - \mathcal{L}[W_{(20+K)T}^{\widetilde{\text{EX}}}\]\|_{\text{TV}} & \leq \frac{7}{500} + \max_{a,b} \mathbb{P}[E_{a,b}^1] + \frac{1}{|V|^2} \sum_{a,b} \mathbb{P}[\sigma_{a,b}(a^*) = a^*] \\ & \quad + \frac{1}{|V|^2} \max \left\{ \frac{23}{25}, \frac{1}{2} + \frac{k-1}{|V|-1} \right\} \sum_{a,b} \mathbb{P}[\sigma_{a,b}(a^*) \neq a^*]. \end{aligned}$$

By assumption, $k \leq |V|/2$ if $|V| < 36$ and $k \in \{3, 4\}$ if $|V| \geq 36$, and so

$$\max \left\{ \frac{23}{25}, \frac{1}{2} + \frac{k-1}{|V|-1} \right\} \leq \frac{33}{34}$$

for all possible combinations of k and $|V|$ being considered here. Combining this bound with that in Assumption 1.1, we obtain:

$$\|\mathcal{L}[U_{(20+K)T}^{\text{EX}}] - \mathcal{L}[W_{(20+K)T}^{\widetilde{\text{EX}}}\]\|_{\text{TV}} \leq \frac{7}{500} + \max_{a,b} \mathbb{P}[E_{a,b}^1] + \frac{1}{5} \left(1 + 4 \cdot \frac{33}{34} \right).$$

But since $\tau_{a,b}$ has the same distribution as $M^{\text{RW}}(I_{10T}(a), \tilde{I}_{10T}(b))$, and G is an easy hypergraph,

$$\max_{a,b} \mathbb{P}[E_{a,b}^1] = \max_{a,b} \mathbb{P}[\tau_{a,b} > KT] \leq \max_{a,b} \mathbb{P}[M^{\text{RW}}(a, b) > KT] \leq \frac{1}{1000}$$

provided $K \geq 10^{10}$. Therefore,

$$\|\mathcal{L}[U_{10^{11}T}^{\text{EX}}] - \mathcal{L}[W_{10^{11}T}^{\widetilde{\text{EX}}}\]\|_{\text{TV}} \leq \frac{8}{500} + \frac{1}{5} \left(1 + 4 \cdot \frac{33}{34} \right) < \frac{497}{500}.$$

Finally, by submultiplicativity of the function

$$\bar{d}(t) := \max_{U,W \in \binom{V}{k}} \|\mathcal{L}[U_t^{\text{EX}}] - \mathcal{L}[W_t^{\widetilde{\text{EX}}}\]\|_{\text{TV}}$$

(see e.g. Lemma 4.12 of [10]), we deduce that

$$\|\mathcal{L}[U_{10^{14} \log(1/\varepsilon)T}^{\text{EX}}] - \mathcal{L}[W_{10^{14} \log(1/\varepsilon)T}^{\widetilde{\text{EX}}}\|\|_{\text{TV}} < \left(\frac{497}{500}\right)^{1000 \log(1/\varepsilon)} < \varepsilon,$$

and so the statement of Lemma 4.4 is proved upon taking $\kappa = 10^{14}$. □

Proof of Lemma 1.10 for easy hypergraphs. We simply apply Lemma 4.4 for the case $|V| \geq 36$ twice, first with $k = 4$ and then with $k = 3$ (and take $\varepsilon = 1/4$ both times). We deduce that

$$T_{\text{EX}(4,f,G)}(1/4) \leq \kappa^2 (\log 4)^2 T_{\text{EX}(2,f,G)}(1/4),$$

and so it suffices to take $\lambda = \kappa^2 (\log 4)^2$. □

4.2 From 4-particle exclusion to 2-particle exclusion: non-easy hypergraphs

We begin with a result showing that for non-easy hypergraphs the average meeting time for two independent random walkers is unlikely to be quick. Intuitively, this follows from the following observations. We know there exists a pair of vertices such that random walkers started from these two states likely take a long time to meet. If we look at where these two walkers are at time of order $T_{\text{RW}(f,G)}(1/4)$, they will be close to uniform. Hence, starting random walkers from a uniform pair we see that they will likely still take a long time to meet. The proofs of Lemmas 4.5 and 4.6, and of Proposition 4.7 are (somewhat technical) extensions of corresponding results of [15], and can be found in Appendix B.

Lemma 4.5. *For every non-easy hypergraph we have*

$$\sum_{\mathbf{u} \in V^2} \frac{\mathbb{P}[M^{\text{RW}}(\mathbf{u}) \leq 20T]}{|V|^2} \leq \frac{1}{1000}.$$

Given a k -tuple $\mathbf{z} \in (V)_k$, we once again write $\mathbf{O}(\mathbf{z}) := \{\mathbf{z}(1), \dots, \mathbf{z}(k)\}$ for the (un-ordered) set of coordinates of \mathbf{z} . For $\mathbf{x} \in V^4$, let \mathbf{x}_t^{RW} be a realisation of $\text{RW}(4, f, G)$ with $\mathbf{x}_0^{\text{RW}} = \mathbf{x}$. Denote by $\Lambda^1, \Lambda^2, \Lambda^3, \Lambda^4$ the (independent) Poisson processes used to generate the edge-ringing times for the four random walkers, and let $\{e_n^1\}_{n \in \mathbb{N}}, \{e_n^2\}_{n \in \mathbb{N}}, \{e_n^3\}_{n \in \mathbb{N}}, \{e_n^4\}_{n \in \mathbb{N}}$ be the four sequences of edge-choices (all as in Section 2.2).

We now define $\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x}))$ to be the first time any two of $\mathbf{x}_t^{\text{RW}}(1), \mathbf{x}_t^{\text{RW}}(2), \mathbf{x}_t^{\text{RW}}(3), \mathbf{x}_t^{\text{RW}}(4)$ first arrive onto the same edge which then rings for one of them. Formally,

$$\begin{aligned} \bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x})) := \inf \left\{ t \geq 0 : \exists 1 \leq i < j \leq 4, e \in \{e_{\Lambda^i[0,t]}^i, e_{\Lambda^j[0,t]}^j\} \right. \\ \left. \text{with } \mathbf{x}_{t-}^{\text{RW}}(i), \mathbf{x}_{t-}^{\text{RW}}(j) \in e \right\}. \end{aligned}$$

Lemma 4.6. *Let $\mathbf{x} \in (V)_4$. Then for any $\varepsilon \in (0, 1)$,*

$$\mathbb{P}[\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x}))_{20T}^{\text{EX}} \leq 20T] \leq 12(\varepsilon + \varepsilon^{-1}2^{-20}) + 25(1 + \varepsilon) \sum_{\mathbf{u} \in V^2} \frac{\mathbb{P}[M^{\text{RW}}(\mathbf{u}) \leq 20T]}{|V|^2}$$

Next, we provide a bound which relates the total-variation distance between two 4-particle exclusion processes to the probability that any two of four independent walkers have ‘met’.

Proposition 4.7. *For any $\mathbf{x} \in (V)_4$ and $s \geq 0$:*

$$\|\mathcal{L}[\mathbf{O}(\mathbf{x}_s^{\text{RW}})] - \mathcal{L}[\mathbf{O}(\mathbf{x}_s^{\text{EX}})]\|_{\text{TV}} \leq \mathbb{P}[\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x})) \leq s].$$

Lemma 4.8. *For every non-easy hypergraph G and any two realisations of $\text{EX}(4, f, G)$, denoted $\{A_t^{\text{EX}}\}$ and $\{B_t^{\text{EX}}\}$, we have*

$$\|\mathcal{L}[A_{40T}^{\text{EX}}] - \mathcal{L}[B_{40T}^{\text{EX}}]\|_{\text{TV}} \leq \mathbb{P} [\bar{M}^{\text{RW}}(A_{20T}^{\text{EX}}) \leq 20T] + \mathbb{P} [\bar{M}^{\text{RW}}(B_{20T}^{\text{EX}}) \leq 20T] + 2^{-18}.$$

Proof. By Proposition 4.7 and the triangle inequality for total-variation, for any $\mathbf{u}, \mathbf{v} \in (V)_4$,

$$\begin{aligned} \|\mathcal{L}[\mathbf{O}(\mathbf{u})_{20T}^{\text{EX}}] - \mathcal{L}[\mathbf{O}(\mathbf{v})_{20T}^{\text{EX}}]\|_{\text{TV}} &\leq \mathbb{P} [\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{u})) \leq 20T] + \mathbb{P} [\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{v})) \leq 20T] \\ &\quad + \|\mathcal{L}[\mathbf{O}(\mathbf{u}_{20T}^{\text{RW}})] - \mathcal{L}[\mathbf{O}(\mathbf{v}_{20T}^{\text{RW}})]\|_{\text{TV}}. \end{aligned} \tag{4.12}$$

An identical argument to that used for equation (4.2) tells us that

$$\|\mathcal{L}[A_{40T}^{\text{EX}}] - \mathcal{L}[B_{40T}^{\text{EX}}]\|_{\text{TV}} \leq \mathbb{E} [\|\mathcal{L}[A_{40T}^{\text{EX}} | A_{20T}^{\text{EX}}] - \mathcal{L}[B_{40T}^{\text{EX}} | B_{20T}^{\text{EX}}]\|_{\text{TV}}].$$

Applying the inequality in (4.12), with any \mathbf{u}, \mathbf{v} satisfying $\mathbf{O}(\mathbf{u}) = A_{20T}^{\text{EX}}$ and $\mathbf{O}(\mathbf{v}) = B_{20T}^{\text{EX}}$, gives

$$\begin{aligned} \|\mathcal{L}[A_{40T}^{\text{EX}}] - \mathcal{L}[B_{40T}^{\text{EX}}]\|_{\text{TV}} &\leq \mathbb{E} [\mathbb{P} [\bar{M}^{\text{RW}}(A_{20T}^{\text{EX}}) \leq 20T | A_{20T}^{\text{EX}}]] \\ &\quad + \mathbb{E} [\mathbb{P} [\bar{M}^{\text{RW}}(B_{20T}^{\text{EX}}) \leq 20T | B_{20T}^{\text{EX}}]] \\ &\quad + \sup_{\mathbf{u}, \mathbf{v} \in (V)_4} \|\mathcal{L}[\mathbf{u}_{20T}^{\text{RW}}] - \mathcal{L}[\mathbf{v}_{20T}^{\text{RW}}]\|_{\text{TV}}. \end{aligned}$$

Using Proposition 2.3 and the contraction principle for the third term on the right-hand side gives the desired result. \square

We are now ready to prove the main result of this subsection.

Proof of Lemma 1.10 for non-easy hypergraphs. We in fact show that for any two realisations of $\text{EX}(4, f, G)$, denoted $\{A_t^{\text{EX}}\}$ and $\{B_t^{\text{EX}}\}$, we have

$$\|\mathcal{L}[A_{40T}^{\text{EX}}] - \mathcal{L}[B_{40T}^{\text{EX}}]\|_{\text{TV}} \leq 1/4.$$

Combining Lemmas 4.6 and 4.8 we have that for every $\varepsilon \in (0, 1)$,

$$\|\mathcal{L}[A_{40T}^{\text{EX}}] - \mathcal{L}[B_{40T}^{\text{EX}}]\|_{\text{TV}} \leq 24(\varepsilon + \varepsilon^{-1}2^{-20}) + 50(1 + \varepsilon) \sum_{\mathbf{u} \in V^2} \frac{\mathbb{P} [M^{\text{RW}}(\mathbf{u}) \leq 20T]}{|V|^2} + 2^{-18}.$$

Now by Lemma 4.5, this becomes

$$\|\mathcal{L}[A_{40T}^{\text{EX}}] - \mathcal{L}[B_{40T}^{\text{EX}}]\|_{\text{TV}} \leq 24(\varepsilon + \varepsilon^{-1}2^{-20}) + \frac{50(1 + \varepsilon)}{1000} + 2^{-18}.$$

Setting $\varepsilon = 10^{-3}$ completes the proof. \square

5 From k -particle exclusion to 2-particle exclusion for small $|V|$

We now prove Lemma 1.7. We begin by showing that any hypergraph G with $|V| < 36$ satisfies

$$\sup_{\mathbf{y} \in V^2} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 10^{10}T_{\text{RW}(f,G)}(1/4)] \leq 1/1000, \tag{5.1}$$

i.e. the hypergraph G is easy. Indeed, by Proposition 2.4, for any $t \geq 2T_{\text{RW}(f,G)}(\varepsilon)$,

$$\sup_{\mathbf{y} \in V^2} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) < t] \geq \frac{(1 - 2\varepsilon)^2}{|V|} \geq \frac{(1 - 2\varepsilon)^2}{36},$$

and so

$$\sup_{\mathbf{y} \in V^2} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 2000T_{\text{RW}(f,G)}(1/4)] \leq 1/1000,$$

which certainly implies (5.1). Since G is easy, we can apply Lemma 4.4 multiple times to deduce that

$$T_{\text{EX}(k,f,G)}(\varepsilon) \leq \kappa^{k-2}(\log(1/4))^{k-3} \log(1/\varepsilon)T_{\text{EX}(2,f,G)}(1/4).$$

However, since $|V| < 36$ and $k \leq |V|/2$ the statement of the proof is complete taking $C = \kappa^{15}(\log(1/4))^{14}$.

6 The chameleon process

Our aim in this section is to construct a continuous-time Markov process which satisfies the properties of $(M_t)_{t \geq 0}$ outlined in Lemma 3.1. We will call this process the *chameleon process*. In Section 7 we will prove Lemma 3.1 by demonstrating that the chameleon process does indeed have the desired properties.

The chameleon process was originally constructed (in a different form but to serve a similar purpose) by Morris ([14]), and then adapted by Oliveira ([15]) to analyse the mixing time of the k -particle interchange process on a graph (as opposed to on a hypergraph, as we consider here). It is built on top of an underlying interchange process, with the aim of helping to describe the distribution of the location of the k th particle in this process, conditional on the locations of the $k - 1$ other particles.

Unlike in a k -particle interchange process which always has k particles, the chameleon process has $|V|$ particles (one at each vertex), although not all particles are distinguishable from each other. In addition, each particle has an associated *colour*: one of black, red, pink and white (which correspond to the processes $\mathbf{z}_t^C, R_t, P_t, W_t$ respectively, appearing in the statement of Lemma 3.1). The movement of particles in the chameleon process follows that of the underlying interchange process in the sense that the locations of particles in both processes are updated using the same functions I as described in the graphical construction of Section 2.2. At some of the updates of the underlying interchange process we will colour some of the red and white particles pink (precisely when this happens is rather involved and is the subject of Section 6.2). To provide some insight into when these *pinkening* events occur, consider the chameleon process of [15]: here, if the vertices at the endpoints of a ringing edge are occupied by a red and a white particle then both of these particles are recoloured pink. In the lazy version of the interchange process on a graph (in which nothing happens with probability 1/2 when an edge rings), when an edge rings with endpoints occupied by a red and a white particle, with probability 1/2 they switch places and with probability 1/2 they do not move. Colouring both particles pink (which should be viewed as half red, half white) encodes the fact that at either vertex just after the edge rings we may have a red particle or a white particle, and these are equally likely.

We wish to use this notion of pinkening to encode similar events in the interchange process on hypergraphs, but the situation here is quite different since more than two particles are moved when an edge rings, and the way in which they move depends on the permutation chosen. As a result, describing precisely when these pinkenings occur for our version of the chameleon process is rather complicated, but the underlying motivation can be explained relatively simply. As in Oliveira’s argument we will use pink particles as a way of tracking particles which are either red or white (equally likely). Whereas Oliveira could make use of laziness to split the conditional distribution among two sites, when dealing with hypergraphs we have to use a new idea of a “twin” permutation. Suppose that an edge e rings and a permutation σ is chosen to move

the particles on that edge. To decide which particles to pinken, we construct a twin permutation $\tilde{\sigma}$ with the property that the trajectories of all black particles in the edge are identical under both permutations (a required property – see part 1 of Lemma 3.1), and such that, viewed marginally, the distribution of $\tilde{\sigma}$ agrees with that of σ . We then look for vertices v such that under σ a red particle is moved to v and under $\tilde{\sigma}$ a white particle is moved to v ; a certain subset of these particles will be pinkened. The simplest example to consider is that of an edge of size 3 which contains one red, one white and one black particle, and for which f_e is constant on S_e . In this case, it is straightforward to construct a twin permutation with these required properties, and the construction is sketched in Figure 4.




Figure 4: Consider an edge of size 3, containing one red, one white and one black particle, and for which f_e is constant on S_e . The six possible permutations are sketched here: in this example the twin of any given permutation σ could be taken to be the permutation immediately above/below σ . Note that in each case the black \bullet particle follows the same trajectory under both σ and its twin; moreover, if σ moves a red/white particle to a vertex v , then its twin moves a white/red particle to v . In this simple example we could therefore pinken the red and white particles, no matter which σ is chosen.

Although this example demonstrates one possibility for generating twin permutations with our desired properties, this is *very* particular to the situation in which f_e is constant on S_e – a much stronger condition than we are imposing in Assumption 1.1. In general, we shall make use of the fact that f_e is constant on conjugacy classes to construct a twin permutation $\tilde{\sigma}$ with the *same cycle structure* as σ . (Note that the twin permutations constructed in Figure 4 do not have the same cycle structure as σ , and so we shall need to come up with an alternative method of pinkening, even when considering edges of size 3.) Figure 5 gives an indication of how $\tilde{\sigma}$ will be produced from knowledge of σ and the particle colours in the case of σ being a single cycle: by modifying the trajectories of four particular particles we are able to ensure that not only does $\tilde{\sigma}$ have the same cycle structure as σ , but that the trajectories of all black particles in the edge are identical under both permutations. It is for this reason (i.e. needing to know the colours of four particular particles) that we are able to relate the mixing time of k particles to that of just four particles in Lemma 1.8.

The chameleon process also updates at additional times (compared to its corresponding interchange). We refer to these additional updates as *depinkings*, as at these times we get the opportunity to collectively recolour all pink particles in the system either red or white. As in [15], we will only perform a depinking once there are a large number of pink particles (compared to the number of red and white) in the system.




Figure 5: If σ is the cycle given by the arrows in the image on the left, then one possible candidate for $\tilde{\sigma}$ is the cycle given by the arrows in the image on the right. Note that all black \bullet particles follow the same trajectories under both cycles, and that the colour of $\tilde{\sigma}(v)$ agrees with that of $\sigma(v)$ for all v .

6.1 Twin permutations

The first step towards constructing the chameleon process is describing how to generate the *twin* permutation $\tilde{\sigma}$ from σ , which is the subject of this section. We begin with the case of σ being either a product of disjoint transpositions or a single cycle (of size at least 3), and then describe how to construct $\tilde{\sigma}$ for a general permutation σ . We conclude the section by describing an algorithm to generate a certain set A which plays a crucial role in the construction of the chameleon process.

6.1.1 Composition of transpositions, and cycles of size at least 3

We begin with some notation: for $d \in \mathbb{N}$ let $[d] = \{1, 2, \dots, d\}$, and let $d' = \lfloor d/4 \rfloor$ (the floor of $d/4$). For convenience we also let $[0] = \{0\}$.

For $d \in 2\mathbb{N}$ we let T_d be the set of products of disjoint transpositions:

$$T_d := \left\{ \prod_{i=1}^{d/2} (a_{2i-1} \ a_{2i}) : 1 \leq a_i \leq d \text{ for all } 1 \leq i \leq d, a_i \neq a_j \text{ for all } i \neq j \right\}.$$

For $\sigma \in T_d$ we define an ordering, denoted \prec , of the transpositions in σ as follows: $(a_i \ a_j) \prec (a_k \ a_\ell)$ if and only if $(a_i \wedge a_j) < (a_k \wedge a_\ell)$. Without loss of generality we shall always suppose that any $\sigma \in T_d$ is written such that

$$(a_1 \ a_2) \prec (a_3 \ a_4) \prec \dots \prec (a_{d-1} \ a_d),$$

and $a_{2i-1} < a_{2i}$ for all $1 \leq i \leq d/2$.

Given a set $A \subseteq [d']$ and a permutation $\sigma \in T_d$, we define the permutation $\beta_A(\sigma)$ to be the result of multiplying σ (on the right) by a particular set of disjoint transpositions, as follows:

$$\beta_A(\sigma) = \sigma \prod_{\substack{i \in A: \\ a_{4i-1} < a_{4i-2}}} (a_{4i-3} \ a_{4i-1})(a_{4i-2} \ a_{4i}). \tag{6.1}$$

The permutation $\beta_A(\sigma)$ satisfies some nice properties, which we put together in the following Lemma. The proofs are straightforward, but it is worth emphasising that part 2 of Lemma 6.1 (that β_A is an involution) holds precisely because in (6.1) we only multiply by transpositions for which $a_{4i-1} < a_{4i-2}$.

Lemma 6.1. *For any $\sigma \in T_d$ and set $A \subseteq [d']$:*

1. $\beta_A(\sigma) \in T_d$;
2. $\beta_A(\beta_A(\sigma)) = \sigma$;

3. for all $x \in [d]$, $\beta_A(\sigma)(x) = \sigma(x)$ unless $x \in \{a_{4i-3}, a_{4i-2}, a_{4i-1}, a_{4i}\}$ for some $i \in A$ with $a_{4i-1} < a_{4i-2}$.

We now move onto cycles of size at least 3. For $d \in \mathbb{N}$ we denote by \mathcal{C}_d the conjugacy class of \mathcal{S}_d (the symmetric group on $[d]$) consisting of cycles of length d . For a cycle $\sigma \in \mathcal{C}_d$ we may write

$$\sigma = (\sigma^0(1) \sigma^1(1) \sigma^2(1) \dots \sigma^{d-1}(1))$$

where for $m \in \mathbb{N}$ we write σ^m for the composition of m copies of σ (and where we may sometimes write $\sigma^d \equiv \sigma^0$ for the identity permutation).

For $d \geq 4$ and $i = 1, \dots, d'$ define the function $\beta_i : [d] \rightarrow [d]$ by

$$\beta_i(j) = \begin{cases} 2d' + 2i - 1 & j = 2i - 1 \\ 2i - 1 & j = 2d' + 2i - 1 \\ j & \text{otherwise.} \end{cases} \quad (6.2)$$

For $d = 3$ we similarly define the function $\beta_0 : [d] \rightarrow [d]$ by

$$\beta_0(j) = \begin{cases} 2 & j = 1 \\ 1 & j = 2 \\ 3 & j = 3. \end{cases} \quad (6.3)$$

For $d \geq 4$ and a cycle $\sigma \in \mathcal{C}_d$, for each $1 \leq i \leq d'$, we define $\beta_i(\sigma) \in \mathcal{C}_d$ to be the permutation satisfying

$$\beta_i(\sigma)^j(1) = \sigma^{\beta_i(j)}(1), \quad j = 1, \dots, d. \quad (6.4)$$

This permutation is clearly a cycle, and may be written as

$$\beta_i(\sigma) = (\sigma^0(1) \dots \sigma^{2i-2}(1) \sigma^{2d'+2i-1}(1) \sigma^{2i}(1) \dots \sigma^{2d'+2i-2}(1) \sigma^{2i-1}(1) \sigma^{2d'+2i}(1) \dots \sigma^{d-1}(1)).$$

For a cycle $\sigma \in \mathcal{C}_3$, we define $\beta_0(\sigma)$ using the same formula as in (6.4), yielding the 3-cycle $\beta_0(\sigma) = \sigma^2 = (1 \sigma^2(1) \sigma(1))$.

Remark 6.2. Note that for the case $d \geq 4$, $\beta_i(\sigma)$ may be obtained from σ by multiplication by the product of two disjoint transpositions:

$$\beta_i(\sigma) = \sigma (\sigma^{2i-2}(1) \sigma^{2d'+2i-2}(1)) (\sigma^{2i-1}(1) \sigma^{2d'+2i-1}(1)).$$

Lemma 6.3. For $d \geq 3$ and any $i, j \in [d']$:

1. The function β_i is self-inverse;
2. functions β_i and β_j commute for $i \neq j$.

Proof. Part 1 follows directly from the definition of β_i . Part 2 only applies when $d \geq 4$, and follows from the observation that the transpositions in Remark 6.2 corresponding to β_i and β_j commute for $i \neq j$. \square

Definition 6.4. Given a set $A \subseteq [d']$ and a cycle $\sigma \in \mathcal{C}_d$, we define $\beta_A : [d] \rightarrow [d]$ to be the composition of the functions $\{\beta_i : i \in A\}$ appearing in (6.2) and (6.3). (Thanks to the second statement of Lemma 6.3 this function is well-defined.) If $A = \emptyset$ then we set β_A to be the identity function.

This in turn defines a cycle $\beta_A(\sigma) \in \mathcal{C}_d$ satisfying

$$\beta_A(\sigma)^j(1) = \sigma^{\beta_A(j)}(1), \quad j = 1, \dots, d. \quad (6.5)$$

When σ is a d -cycle the permutation $\beta_A(\sigma)$ satisfies analogous properties to those already observed to hold (in Lemma 6.1) for $\beta_A(\sigma)$ when $\sigma \in T_d$. The proofs all follow simply from the definition of β_i and Lemma 6.3. For each $i \in [d']$, write

$$H_i = \begin{cases} \{2i - 2, 2i - 1, 2d' + 2i - 2, 2d' + 2i - 1\} & \text{if } d' \geq 1 \text{ (i.e. } d \geq 4), \\ \{1, 2, 3\} & \text{if } d' = 0 \text{ (i.e. } d = 3), \end{cases} \quad (6.6)$$

and for $A \subseteq [d']$ let $H_A = \bigcup_{i \in A} H_i$.

Lemma 6.5. *For any $d \geq 3$, $\sigma \in \mathcal{C}_d$ and set $A \subseteq [d']$:*

1. $\beta_A(\sigma) \in \mathcal{C}_d$;
2. $\beta_A(\beta_A(\sigma)) = \sigma$;
3. for all $x \in [d]$, $\beta_A(\sigma)(x) = \sigma(x)$ unless $x = \sigma^j(1)$ for some $j \in H_A$.

So far we have defined a method for producing a permutation $\beta_A(\sigma)$ in the event that σ is either a product of disjoint transpositions or a cycle of length at least 3. Now consider what happens when we apply the function β_A to a permutation chosen uniformly from either \mathcal{C}_d or T_d (for some d). Clearly, for any set A chosen independently of σ , the resulting permutation $\beta_A(\sigma)$ will be uniformly distributed on the same conjugacy class as σ . Most importantly, this remains true even when A is allowed to depend upon σ , as long as a certain condition is met, as explained in the following Lemma. We denote by \mathcal{P}^Ω the power set of a set Ω .

Lemma 6.6. *Let G_d denote either of the conjugacy classes \mathcal{C}_d ($d \geq 3$) or T_d ($d \in 2\mathbb{N}$). Suppose that $A : G_d \rightarrow \mathcal{P}^{[d']}$ satisfies for all $\sigma \in G_d$,*

$$A(\beta_{A(\sigma)}(\sigma)) = A(\sigma), \quad (6.7)$$

and that σ is chosen uniformly from G_d . Then $\beta_{A(\sigma)}(\sigma)$ is also uniform on G_d . Moreover, if we average over the input permutation σ , then the output $\beta_{A(\sigma)}(\sigma)$ is independent of the choice of A .

Proof. Given the permutation σ , let $\tilde{\sigma} = \beta_{A(\sigma)}(\sigma)$. The assumption on A says that $A(\tilde{\sigma}) = A(\sigma)$. Since β_A is an involution (Lemmas 6.1 and 6.5) it follows that

$$\beta_{A(\tilde{\sigma})}(\tilde{\sigma}) = \beta_{A(\sigma)}(\tilde{\sigma}) = \beta_{A(\sigma)}(\beta_{A(\sigma)}(\sigma)) = \sigma.$$

Thus the function $\beta_{A(\cdot)}(\cdot)$ is self-inverse. □

Although Lemma 6.6 is relatively simple, its importance should be emphasised at this point. We shall make use of the function β_A to generate the random permutations $\tilde{\sigma}$ used in the construction of the chameleon process, and in doing so the input A will depend on the state of the chameleon process. The second part of Lemma 6.6 will be used to guarantee that the permutation $\tilde{\sigma} = \beta_{A(\sigma)}(\sigma)$ is independent of A . (The permutations $\tilde{\sigma}$ will be used to generate an interchange process $\tilde{\mathbf{x}}^{\text{IP}}$, and so it will be crucial that these do not depend on the state of the process.)

6.1.2 General permutations

By combining the ideas from the previous two sections we can now describe the algorithm for the construction of the twin permutation $\tilde{\sigma}$ (which will be given by $\beta_{A(\sigma)}(\sigma)$ for some function $\beta_{A(\cdot)}(\cdot)$ to be defined) when $\sigma \in \mathcal{S}_n$ is a general permutation. The first step is to decompose the input permutation σ into its canonical cyclic decomposition form. Indeed,

except for transpositions, the function $\beta_{A(\cdot)}(\cdot)$ will act independently on each cycle in a given permutation's decomposition.

Suppose σ has canonical cyclic decomposition form (where we omit fixed points):

$$\sigma = \rho_0 \circ \rho_1 \circ \cdots \circ \rho_K, \tag{6.8}$$

where K denotes the number of cycles in σ of size at least 3, and ρ_0 is a (possibly empty) product of disjoint transpositions.

For $i = 0, 1, \dots, K$ we write m_i for the minimal element of ρ_i , and write d_i for the size of the non-trivial orbit of ρ_i . (For example, if $\sigma = (14)(29)(37685) \in \mathcal{S}_9$ then $K = 1$, $m_0 = 1$, $m_1 = 3$, $d_0 = 4$ and $d_1 = 5$.) Given the elements of the non-trivial orbit of ρ_i , there is an obvious natural bijection between permutations of those elements and permutations of the set $[d_i] = \{1, \dots, d_i\}$, in which the minimal element m_i is mapped to 1. Rather than writing out this correspondence in detail, in order to ease notation in what follows we shall simply consider ρ_i to be a member of the set \mathcal{C}_{d_i} etc, even though the set of elements belonging to ρ_i will not in general be $\{1, \dots, d_i\}$.

With this understanding in mind, suppose that $A(\sigma)$ is a vector of the form

$$A(\sigma) = (A_0(\rho_0), A_1(\rho_1), \dots, A_K(\rho_K)), \tag{6.9}$$

where $A_0 : T_{d_0} \rightarrow \mathcal{P}^{[d_0]}$ and $A_i : \mathcal{C}_{d_i} \rightarrow \mathcal{P}^{[d_i]}$ for $i = 1, \dots, K$. Then we can easily extend the idea of our functions β_A to apply to general permutations.

Definition 6.7. Let $\sigma \in \mathcal{S}_n$ be a permutation with cyclic decomposition (6.8), and assume that A is a function on \mathcal{S}_n satisfying (6.9). Then we define $\tilde{\beta}_{A(\sigma)}(\sigma)$ to be the composition of the permutations obtained by applying the functions $\beta_{A_i(\rho_i)}$ separately to each ρ_i :

$$\tilde{\beta}_{A(\sigma)}(\sigma) = \prod_{i=0}^K \beta_{A_i(\rho_i)}(\rho_i),$$

where $\beta_{A_i(\rho_i)}(\rho_i)$ are as defined in Section 6.1.1 (but with m_i replacing the element 1, as already explained).

Definition 6.7 says that $\tilde{\beta}_{A(\sigma)}(\sigma)$ is obtained from σ by modifying each of its cycles of size at least 3, and the set of disjoint transpositions, independently using functions $\beta_{A_i(\cdot)}(\cdot)$ with which we are already familiar. We therefore have the following corollary to Lemmas 6.1 and 6.5.

Corollary 6.8. For any $\sigma \in \mathcal{S}_n$ and function A on \mathcal{S}_n satisfying (6.9):

1. $\tilde{\beta}_{A(\sigma)}(\sigma)$ belongs to the same conjugacy class as σ ;
2. $\tilde{\beta}_{A(\sigma)}(\tilde{\beta}_{A(\sigma)}(\sigma)) = \sigma$;
3. for all $x \in [n]$, $\tilde{\beta}_{A(\sigma)}(x) = \sigma(x)$ unless $x \in \{a_{4i-3}, a_{4i-2}, a_{4i-1}, a_{4i}\}$ for some $i \in A_0(\rho_0)$ with $a_{4i-1} < a_{4i-2}$, or $x = \sigma^j(m_i)$ for some $j \in \cup_{i=1}^K H_{A_i(\rho_i)}$.

Furthermore, note that if we choose a random permutation $\sigma \in \mathcal{S}_n$ according to a law f which is constant on conjugacy classes, then given the sizes of the cycles in the decomposition of σ , the elements of $[n]$ belonging to each cycle are (marginally) uniform. We can therefore also obtain a corollary to Lemma 6.6:

Corollary 6.9. Suppose that A is a function on \mathcal{S}_n satisfying (6.9), and that for all $\sigma \in \mathcal{S}_n$ with cyclic decomposition (6.8) and each $i = 0, 1, \dots, K$,

$$A_i(\beta_{A_i(\rho_i)}(\rho_i)) = A_i(\rho_i).$$

If σ is chosen according to law f on \mathcal{S}_n which is constant on conjugacy classes then $\tilde{\beta}_{A(\sigma)}(\sigma)$ also has law f on \mathcal{S}_n . Moreover, if we average over the input permutation σ , then the output $\tilde{\beta}_{A(\sigma)}(\sigma)$ is independent of the choice of A .

6.1.3 Choosing the set A

We have described in Section 6.1.2 how to generate the twin permutation $\tilde{\sigma} = \tilde{\beta}_{A(\sigma)}(\sigma)$ from a permutation σ such that it has the same law as σ . We now detail our method for choosing the vector $A(\sigma)$ appearing in the definition of $\tilde{\beta}_A$, in such a way that the conditions of Corollary 6.9 are satisfied; an illustrative example can be found in Figures 6 and 7. Our choice of A will depend not only on σ but also on particular subsets of vertices in the edge under consideration. (Later on these subsets will be specified in the chameleon process, but for now we keep them as general subsets.) Indeed, given an edge $e \in E$ and a permutation $\sigma \in \mathcal{S}_e$, the function A is of the form $A(R, W, \sigma)$, where R and W are two disjoint subsets of V .

Recall the definition of the set H_i in (6.6), and the canonical cyclic decomposition of σ from (6.8) in which K denotes the number of cycles of size at least 3 in σ . For a set of integers B let us write $\sigma^B(x) = \{\sigma^i(x) : i \in B\}$. Then for each $1 \leq i \leq K$, we define

$$A_i(R, W, \rho_i) = \begin{cases} \left\{ \begin{array}{l} \left\{ j \in [d'_i] : \rho_i^{H_j}(m_i) \in \{\{r_1, w_1, w_2\}, \{w_1, r_1, r_2\}\} \right. \\ \left. \text{for some } r_1, r_2 \in R, w_1, w_2 \in W \right\} \end{array} \right\} & \text{if } d_i = 3 \\ \left\{ \begin{array}{l} \left\{ j \in [d'_i] : \rho_i^{H_j}(m_i) \in \{\{r_1, w_1, w_2, w_3\}, \{w_1, r_1, r_2, r_3\}\} \right. \\ \left. \text{for some } r_1, r_2, r_3 \in R, w_1, w_2, w_3 \in W \right\} \end{array} \right\} & \text{if } d_i \geq 4. \end{cases} \tag{6.10}$$

Recall that ρ_0 denotes the composition of all disjoint transpositions in σ . Using our usual ordering we can write

$$\rho_0 = \prod_{i=1}^{d_0/2} (a_{2i-1} a_{2i})$$

for some $d_0 \in 2\mathbb{N}$ and elements $a_j \in e$, where $a_{2i-1} < a_{2i}$ for all $1 \leq i \leq d_0/2$. This allows us to define

$$A_0(R, W, \rho_0) = \left\{ \begin{array}{l} \left\{ j \in [d'_0] : \{a_{4j-3}, a_{4j-2}, a_{4j-1}, a_{4j}\} \in \{\{r_1, w_1, w_2, w_3\}, \{w_1, r_1, r_2, r_3\}\} \right. \\ \left. \text{for some } r_1, r_2, r_3 \in R, w_1, w_2, w_3 \in W \right\} \end{array} \right\}. \tag{6.11}$$

Finally, we define $A(R, W, \sigma)$ to be the vector

$$A(R, W, \sigma) = (A_0(R, W, \rho_0), A_1(R, W, \rho_1), A_2(R, W, \rho_2), \dots, A_K(R, W, \rho_K)). \tag{6.12}$$

We conclude this section by showing that this construction of A satisfies the conditions of Corollary 6.9. This in turn guarantees that the permutation $\tilde{\sigma} = \tilde{\beta}_{A(\sigma)}(\sigma)$ has the same law on \mathcal{S}_e as σ .

Lemma 6.10. *Fix an edge $e \in E$. For any disjoint subsets R, W of V and permutation $\sigma \in \mathcal{S}_e$, the functions $A_i(R, W, \rho_i)$ defined in (6.10) and (6.11) satisfy*

$$A_i(R, W, \beta_{A_i(R, W, \rho_i)}(\rho_i)) = A_i(R, W, \rho_i).$$

Proof. The idea here is that, as in part 3 of Corollary 6.8, $\beta_{A_i(R, W, \rho_i)}(\rho_i)(x) = \rho_i(x)$ unless x belongs to a special set of three or four elements (whose exact definition depends upon the conjugacy class of ρ_i). Furthermore, $\beta_{A_i(R, W, \rho_i)}$ permutes all elements of such a special set amongst themselves, and so the numbers of red and white vertices within the set are unchanged by the action of $\beta_{A_i(R, W, \rho_i)}$. (See Figure 7 again for a pictorial example.)




Figure 6: An example of a permutation applied to an edge of size 25. Particles at vertices in $R = \{1, 2, 5, 6, 9, 12, 20, 22\}$ are coloured red \bullet and at vertices in $W = \{3, 4, 8, 10, 14, 16, 19, 21, 24\}$ are coloured white \circ . (The other colours – black \bullet and pink \circ – will become important later, but may be ignored for now.) Arrows are used to denote the input permutation σ so here we see that $\sigma = (5\ 21)(8\ 10)(16\ 20)(3\ 12\ 22)(1\ 6\ 17\ 18\ 19\ 2\ 13\ 25\ 24\ 9\ 7\ 15\ 14\ 4\ 11\ 23)$. Since the first pair of transpositions $(5\ 21)(8\ 10)$ involve one red and three white particles, we deduce that $A_0 = \{1\}$; similarly, $A_1 = \{0\}$. Looking at the large cycle ρ_2 , we see that both of the sets $\rho_2^{H_1}(1) = \{1, 6, 24, 9\}$ and $\rho_2^{H_3}(1) = \{19, 2, 14, 4\}$ contain a 3:1 split of reds:whites or whites:reds, and so $A_2 = \{1, 3\}$. Putting these together we arrive at $A = (\{1\}, \{0\}, \{1, 3\})$.




Figure 7: A pictorial description of the result of using the set A derived in Figure 6 to construct the twin permutation $\tilde{\sigma} = \tilde{\beta}_{A(R,W,\sigma)}(\sigma)$; here we see that $\tilde{\sigma} = (5\ 10)(8\ 21)(16\ 20)(3\ 22\ 12)(1\ 9\ 17\ 18\ 19\ 4\ 13\ 25\ 24\ 6\ 7\ 15\ 14\ 2\ 11\ 23)$.

We provide some details here for the case when ρ_i is a cycle of length $d_i \geq 4$: the arguments for 3-cycles and ρ_0 are similar. Suppose $j \in A_i(R, W, \rho_i)$. Without loss of

generality, suppose that

$$\begin{aligned} \rho_i^{2j-2}(m_i) \in W, \quad \rho_i^{2d'_i+2j-2}(m_i) \in R, \\ \rho_i^{2j-1}(m_i) \in W, \quad \rho_i^{2d'_i+2j-1}(m_i) \in W. \end{aligned}$$

Since $j \in A_i(R, W, \rho_i)$, from equation (6.5) we deduce that

$$\begin{aligned} \beta_{A_i(R, W, \rho_i)}(\rho_i)^{2d'_i+2j-2}(m_i) &= \rho_i^{2j-2}(m_i) \in W, \\ \beta_{A_i(R, W, \rho_i)}(\rho_i)^{2j-2}(m_i) &= \rho_i^{2d'_i+2j-2}(m_i) \in R, \\ \beta_{A_i(R, W, \rho_i)}(\rho_i)^{2d'_i+2j-1}(m_i) &= \rho_i^{2j-1}(m_i) \in W, \\ \beta_{A_i(R, W, \rho_i)}(\rho_i)^{2j-1}(m_i) &= \rho_i^{2d'_i+2j-1}(m_i) \in W. \end{aligned}$$

Therefore $k \in A_i(R, W, \beta_{A_i(R, W, \rho_i)}(\rho_i))$. The other cases follow similarly. This shows that $A_i(R, W, \rho_i) \subseteq A_i(R, W, \beta_{A_i(R, W, \rho_i)}(\rho_i))$, but an identical argument shows the reverse implication and we deduce the result. \square

6.2 Construction of the chameleon process

In this section we detail the construction of the chameleon process. The connection to the algorithm described in the previous section to generate $\tilde{\sigma}$ from σ will be made clear in Lemma 6.12. In order to deal with edges of size 2, it will be convenient to modify the graphical construction of $\text{IP}(k, f, G)$ introduced in Section 2.2, by doubling the rate at which edges ring, and compensating for this by making the process lazy.

More formally, consider the following ingredients:

1. a Poisson process Λ of rate $2|E|$;
2. an i.i.d. sequence of E -valued random variables $\{e_n\}_{n \in \mathbb{N}}$;
3. an i.i.d. sequence of permutations $\{\sigma_n\}_{n \in \mathbb{N}}$ with $\sigma_n \in \mathcal{S}_{e_n}$ for each $n \in \mathbb{N}$ and with $\mathbb{P}[\sigma_n = \sigma] = f_{e_n}(\sigma)$;
4. an i.i.d. sequence of coin flips $\{\theta_n\}_{n \in \mathbb{N}}$ with $\mathbb{P}[\theta_n = 1] = \mathbb{P}[\theta_n = 0] = 1/2$.

We now define $\sigma_n^{\theta_n}$ to equal σ_n if $\theta_n = 1$ and to be the identity if $\theta_n = 0$. We modify the definition of the maps $I_{[s,t]}$ from Section 2.2 as follows:

$$I_{[s,t]} = \sigma_{e_{\Lambda[0,t]}}^{\theta_{\Lambda[0,t]}} \circ \sigma_{e_{\Lambda[0,t]-1}}^{\theta_{\Lambda[0,t]-1}} \circ \dots \circ \sigma_{e_{\Lambda[0,s]+1}}^{\theta_{\Lambda[0,s]+1}}.$$

The thinning property of Poisson processes implies that the joint distribution of the maps $I_{[s,t]}$, $0 \leq s \leq t < \infty$, is the same as in Section 2.2. The chameleon process will be built on top of this modified interchange process.

6.2.1 Formal description of the chameleon process

Given a $(k-1)$ -tuple $\mathbf{z} \in (V)_{k-1}$, recall that $\mathbf{O}(\mathbf{z}) := \{\mathbf{z}(1), \dots, \mathbf{z}(k-1)\}$ denotes the (unordered) set of coordinates of \mathbf{z} . The chameleon process is a continuous-time Markov process with state-space

$$\Omega_k(V) := \{(\mathbf{z}, R, P, W) : \mathbf{z} \in (V)_{k-1}, \text{ and sets } \mathbf{O}(\mathbf{z}), R, P, W \text{ partition } V\}.$$

A particle at vertex v is said to be *red* if $v \in R$, *white* if $v \in W$, *pink* if $v \in P$ and *black* if $v \in \mathbf{O}(\mathbf{z})$. Note that, due to the nature of the state-space, we can distinguish between the

various black particles, whereas any two red/white/pink particles are indistinguishable from each other.

We denote the state at time t of the chameleon process started from $M_0 = (\mathbf{z}, R, P, W)$ as

$$M_t = (\mathbf{z}_t, R_t, P_t, W_t).$$

We say that a particle at vertex v at time t is *red at time t* if $v \in R_t$ (and similarly for the other colours). We define now a notion of *ink*, which represents the amount of *redness* either at a vertex or in the whole system. A vertex v has 1 unit of ink at time t if $v \in R_t$ and half a unit if $v \in P_t$. Formally then, we define for each $v \in V$ and $t \geq 0$,

$$\text{ink}_t(v) := \mathbf{1}_{\{v \in R_t\}} + \frac{1}{2} \mathbf{1}_{\{v \in P_t\}}. \tag{6.13}$$

We are now able to complete our formal definition of the chameleon process corresponding to an interchange process on a hypergraph. We set $T = 20T_{\text{EX}(4,f,G)}$, and call T the *phase length*. As stated previously, the chameleon process is time-inhomogeneous, and behaves differently depending on which phase we are in. There will be just two different kinds of phase: those in which no colour-changing is permitted and particles are just moved around the graph according to the underlying interchange process; and those in which colour-changing (pinkening of red and white particles) can occur. Furthermore, there will be (deterministic) times at which *depinking* can occur. To be more precise, the chameleon process is updated at the incident times $\{\tau_n\}$ of the Poisson process Λ and also at deterministic times $2iT$, $i \in \mathbb{N}$.

To describe which particles are pinkened during a colour-changing phase, let $\sigma (= \sigma_n)$ be the permutation applied to some edge $e (= e_n)$ at time $t = \tau_n$ and once again recall the cyclic decomposition from (6.8):

$$\sigma = \rho_0 \circ \rho_1 \circ \dots \circ \rho_K.$$

Given that t is in a colour-changing phase, we define subsets of V in the following way.

For cycles ρ_i with $d_i = 3$, recall that $A_i(R_{t-}, W_{t-}, \rho_i)$ is either equal to $\{0\}$ or \emptyset . For $j \in A_i(R_{t-}, W_{t-}, \rho_i)$ we define

$$L_t^{i,j} := \begin{cases} \{r, \rho_i(r)\} & \text{if } \rho_i^{\{0,1,2\}}(m_i) = \{r, w_1, w_2\} \text{ for some } r \in R_{t-}, w_1, w_2 \in W_{t-}, \\ \{w, \rho_i(w)\} & \text{if } \rho_i^{\{0,1,2\}}(m_i) = \{w, r_1, r_2\} \text{ for some } w \in W_{t-}, r_1, r_2 \in R_{t-}. \end{cases}$$

We note that, by construction, if $j = 0$ then $L_t^{i,j}$ contains two vertices, with one in R_{t-} and the other in W_{t-} .

For cycles ρ_i with $d_i \geq 4$, we define a set $L_t^{i,j}$ for each $j \in A_i(R_{t-}, W_{t-}, \rho_i)$ as follows:

- if $|\rho_i^{H_j}(m_i) \cap R_{t-}| = 1$ (so one vertex in the set $\rho_i^{H_j}(m_i)$ contains a red particle, and the other three contain white particles), then set

$$L_t^{i,j} := \begin{cases} \{\rho_i^{2j-2}(m_i), \rho_i^{2d'_i+2j-2}(m_i)\} & \text{if } |\{\rho_i^{2j-2}(m_i), \rho_i^{2d'_i+2j-2}(m_i)\} \cap R_{t-}| = 1, \\ \{\rho_i^{2j-1}(m_i), \rho_i^{2d'_i+2j-1}(m_i)\} & \text{otherwise.} \end{cases}$$

- alternatively, if $|\rho_i^{H_j}(m_i) \cap W_{t-}| = 1$ (so one vertex in the set $\rho_i^{H_j}(m_i)$ contains a white particle, and the other three contain red particles), then set

$$L_t^{i,j} := \begin{cases} \{\rho_i^{2j-2}(m_i), \rho_i^{2d'_i+2j-2}(m_i)\} & \text{if } |\{\rho_i^{2j-2}(m_i), \rho_i^{2d'_i+2j-2}(m_i)\} \cap W_{t-}| = 1, \\ \{\rho_i^{2j-1}(m_i), \rho_i^{2d'_i+2j-1}(m_i)\} & \text{otherwise.} \end{cases}$$

Once again, this ensures that $L_t^{i,j}$ contains two vertices, one in R_{t-} and the other in W_{t-} .

For ρ_0 (the product of disjoint transpositions), we proceed similarly. For each $j \in A_0(R_{t-}, W_{t-}, \rho_0)$ satisfying $a_{4j-1} < a_{4j-2}$, we define $L_t^{0,j}$ as follows:

- if $|\{a_{4j-3}, a_{4j-2}, a_{4j-1}, a_{4j}\} \cap R_{t-}| = 1$, then set

$$L_t^{0,j} := \begin{cases} \{a_{4j-3}, a_{4j-1}\} & \text{if } |\{a_{4j-3}, a_{4j-1}\} \cap R_{t-}| = 1, \\ \{a_{4j-2}, a_{4j}\} & \text{otherwise;} \end{cases}$$

- alternatively, if $|\{a_{4j-3}, a_{4j-2}, a_{4j-1}, a_{4j}\} \cap W_{t-}| = 1$, then set

$$L_t^{0,j} := \begin{cases} \{a_{4j-3}, a_{4j-1}\} & \text{if } |\{a_{4j-3}, a_{4j-1}\} \cap W_{t-}| = 1, \\ \{a_{4j-2}, a_{4j}\} & \text{otherwise.} \end{cases}$$

(If $a_{4j-1} > a_{4j-2}$ then set $L_t^{0,j} = \emptyset$.)

We then let

$$L_t := \bigcup_{i=0}^K \bigcup_{j \in A_i(R_{t-}, W_{t-}, \rho_i)} L_t^{i,j}.$$

Recall the example in Figure 6 (and suppose $R = R_{t-}$ and $W = W_{t-}$). Here we obtain $L_t^{0,1} = \{5, 8\}$, $L_t^{1,0} = \{3, 12\}$, $L_t^{2,1} = \{1, 24\}$ and $L_t^{2,3} = \{2, 4\}$, and hence $L_t = \{1, 2, 3, 4, 5, 8, 12, 24\}$.

The particles at the pairs of vertices selected in this way are those that we wish to pinken at time t . However, it turns out to be useful to limit the number of pinkenings that can occur (during a single colour-changing phase) so that the total number of pinks cannot exceed either the number of reds or the number of whites (this will be crucial to be able to appeal directly to a result of [15] in the proof of Lemma 7.2). In order to achieve this, we pick (arbitrarily) a subset L_t^* of

$$\left\{ L_t^{i,j} : i = 0, 1, \dots, K, j \in A_i(R_{t-}, W_{t-}, \rho_i) \right\}$$

with the property that $|L_t^*|$ is as large as possible while still satisfying

$$|P_{t-}| + 2|L_t^*| \leq \min\{|R_t|, |W_t|\} = \min\{|R_{t-}|, |W_{t-}|\} - |L_t^*|,$$

i.e.

$$|L_t^*| \leq \frac{1}{3} (\min\{|R_{t-}|, |W_{t-}|\} - |P_{t-}|).$$

Note that L_t^* is a set of pairs of vertices, with each pair containing one red and one white particle. Finally, we let \bar{L}_t be the union of the elements of L_t^* .

It is precisely the particles at vertices in \bar{L}_t that we will pinken at time t .

Box 6.11. Formal description of chameleon process updates

There are three kinds of updates to the chameleon process – which update is performed at time t depends on the value of t .

Constant-colour phases: For $t = \tau_n \in (2(i-1)T, (2i-1)T]$, update as the interchange process:

$$(z_t, R_t, P_t, W_t) = (\sigma_n^{\theta_n}(z_{t-}), \sigma_n^{\theta_n}(R_{t-}), \sigma_n^{\theta_n}(P_{t-}), \sigma_n^{\theta_n}(W_{t-})).$$

Colour-changing phases: For $t = \tau_n \in ((2i-1)T, 2iT]$, update as interchange (i.e. update as in a constant-colour phase) unless $|P_{t-}| < \min\{|R_{t-}|, |W_{t-}|\}$ and we are in one of the following two situations:

- $|e_n| > 2$, $\theta_n = 1$ and $\sigma_n \neq id$.

In this case, we pinken all particles in the set \bar{L}_t (half of which belong to R_{t-} , the others to W_{t-} , by design).

Regardless of which particles are pinkened, we then update as the interchange process at this time (using σ_n). Formally then, we update as

$$(\mathbf{z}_t, R_t, P_t, W_t) = (\sigma_n(\mathbf{z}_{t-}), \sigma_n(R_{t-} \setminus (\bar{L}_t \cap R_{t-})), \sigma_n(P_{t-} \cup \bar{L}_t), \sigma_n(W_{t-} \setminus (\bar{L}_t \cap W_{t-}))).$$

- $e_n = \{w, r\}$ for some $w \in W_{t-}$, $r \in R_{t-}$, and $\sigma_n \neq id$.

In this case, we pinken both particles on the edge. Formally update as

$$(\mathbf{z}_t, R_t, P_t, W_t) = (\mathbf{z}_{t-}, R_{t-} \setminus \{r\}, P_{t-} \cup \{r, w\}, W_{t-} \setminus \{w\}).$$

Depinking: For $t = 2iT$ with $i \in \mathbb{N}$, if $|P_{t-}| \geq \min\{|R_{t-}|, |W_{t-}|\}$ then we generate a Bernoulli(1/2) random variable Y_i : if $Y_i = 1$, we colour all pink particles red, otherwise we colour all pink particles white. Hence the update is

$$(\mathbf{z}_t, R_t, P_t, W_t) = \begin{cases} (\mathbf{z}_{t-}, R_{t-} \cup P_{t-}, \emptyset, W_{t-}) & \text{if } Y_i = 1, \\ (\mathbf{z}_{t-}, R_{t-}, \emptyset, W_{t-} \cup P_{t-}) & \text{if } Y_i = 0. \end{cases}$$

Recall again the example from Figure 6, and suppose this represents the state at time $t-$ of a chameleon process. Then Figure 8 represents the state at time t (assuming that t belongs to a colour-changing phase, and that the associated random variable θ equals 1).




Figure 8: The result of updating the chameleon process from the state pictured in Figure 6. Particles belonging to the set $L_t = \{1, 2, 3, 4, 5, 8, 12, 24\}$ have been pinkened, and then all particles have been moved according to the permutation $\sigma = (5\ 21)(8\ 10)(16\ 20)(3\ 12\ 22)(1\ 6\ 17\ 18\ 19\ 2\ 13\ 25\ 24\ 9\ 7\ 15\ 14\ 4\ 11\ 23)$. (The number of particles that we are allowed to pinken depends upon the values of $|R_{t-}|$ and $|W_{t-}|$ of course, but here we have assumed for simplicity that $\bar{L}_t = L_t$.)

The connection between the permutations $\tilde{\beta}_A(\sigma_n)$, which we spent time developing in Section 6.1, and the chameleon process is made explicit in the following lemma, which shall later be employed in the proof of part 2 of Lemma 3.1.

Lemma 6.12. *Suppose $t = \tau_n$ is in a colour-changing phase and $|e_n| > 2$. Then there exists a permutation $g : V \rightarrow V$ such that both of the following statements hold:*

1. *for any vertex u containing (at time $t-$) a particle which is pinkened at time t in the chameleon process,*

$$u \in R_{t-} \quad \text{iff} \quad g(u) \in W_{t-} \quad \text{and} \quad u \in W_{t-} \quad \text{iff} \quad g(u) \in R_{t-};$$

2. *for any vertex u containing a particle which is not pinkened at time t , the particle at vertex $g(u)$ at time $t-$ has the same colour at time t as the particle at vertex u at time $t-$.*

Moreover, we can take g to satisfy

$$\tilde{\beta}_{A(R_{t-}, W_{t-}, \sigma_n)}(\sigma_n)(g(u)) = \sigma_n(u).$$

Proof. This follows simply by comparing the construction of $\tilde{\beta}_{A(R_{t-}, W_{t-}, \sigma_n)}(\sigma_n)$ with the construction of the chameleon process. □

7 Properties of the chameleon process

In this section we show that the chameleon process satisfies the properties outlined in Lemma 3.1. Part 1 follows immediately from the construction of the chameleon process, since each black particle moves identically in the chameleon process and the underlying interchange process.

In order to prove the other three parts, we will need to understand the evolution of the total amount of ink in the chameleon process. We first of all note that the number of pink particles accumulates over time until we have a large number of them; at the next depinking time all pink particles are recoloured (either red or white) and the process of accumulation starts again. The process will continue in this manner until either we have no white particles or we have no red particles (which will occur immediately after some depinking). At this point, no more pink particles can be made and so there is no more recolouring of particles. In order to bound the mixing time of the interchange process we need a good understanding of how quickly the chameleon process reaches the state where no more recolouring can occur. There are two factors which affect this: the time we must wait between depinking events and how the process behaves at depinking times.

Writing $\mathbf{x} = (\mathbf{z}, x)$, for each $j \in \mathbb{N}$ let $D_j(\mathbf{x})$ denote the j th depinking time of a chameleon process started from state $(\mathbf{z}, \{x\}, \emptyset, V \setminus (\mathbf{O}(\mathbf{z}) \cup \{x\})) \in \Omega_k(V)$. Let $\text{ink}_t^{\mathbf{x}}$ denote the total amount of ink in the process at time t ; note that $0 \leq \text{ink}_t^{\mathbf{x}} \leq |V| - k + 1$. Motivated by [15], recall that in part 2 of Lemma 3.1 we defined the event

$$\text{Fill}^{\mathbf{x}} := \left\{ \lim_{t \rightarrow \infty} \text{ink}_t^{\mathbf{x}} = |V| - k + 1 \right\}.$$

This is the event that all initially-white particles are eventually coloured red. We shall make use of the following result concerning $\text{ink}_t^{\mathbf{x}}$, whose proof may be found in [15]. It is applicable in this setting because the event $\text{Fill}^{\mathbf{x}}$ is independent of \mathbf{z}_t^C (as it depends only on the outcomes of coin-flips at depinking times and these do not affect \mathbf{z}_t^C) and because $\text{ink}_t^{\mathbf{x}}$ is a martingale (clear from the construction). We note also that this result is identical to Lemma 3.1 part 4, and thus serves as its proof.

Proposition 7.1. Fix $\mathbf{x} = (\mathbf{z}, x) \in (V)_k$. For each $\mathbf{c} \in (V)_{k-1}$ and $t \geq 0$,

$$\mathbb{P} [\{\mathbf{z}_t^C = \mathbf{c}\} \cap \text{Fill}^{\mathbf{x}}] = \frac{\mathbb{P} [\mathbf{z}_t^C = \mathbf{c}]}{|V| - k + 1}.$$

Consider now the expectation on the right-hand side of the statement of Lemma 3.2: an identical argument to that in Section 6 of [15] shows that this can be bounded in terms of the tail probability of the time of the j th depinking.

Lemma 7.2. There exist positive constants c_1 and c_2 such that for every $j \in \mathbb{N}$,

$$\sup_{\mathbf{x} \in (V)_k} \mathbb{E} \left[1 - \frac{\text{ink}_t^{\mathbf{x}}}{|V| - k + 1} \mid \text{Fill}^{\mathbf{x}} \right] \leq c_1 \sqrt{|V|} e^{-c_2 j} + \sup_{\mathbf{x} \in (V)_k} \mathbb{P} [D_j(\mathbf{x}) > t \mid \text{Fill}^{\mathbf{x}}].$$

We therefore see that we need good control on the probability that there have only been a few depinkings by time t . Here we cannot simply rely on results from [15], since our chameleon process constructed in Section 6 clearly obeys very different dynamics. We shall need the following fundamental result – a lower bound on the number of red particles that are lost (due to pinkening) during a colour-changing phase of the chameleon process (where we start the phase with more white particles than red). The proof is deferred to Section 7.1.

Lemma 7.3. Suppose $|V| \geq 36$ and consider a chameleon process with initial configuration (\mathbf{z}, R, P, W) satisfying $|P| < |R| \leq |W|$. Then

$$\mathbb{E}[|R_{2T-}|] \leq (1 - 10^{-6})|R|.$$

We use this to bound the probability appearing in the statement of Lemma 7.2.

Lemma 7.4. There exists a universal constant $\kappa_1 > 0$ such that for every interchange process on a regular hypergraph $G = (V, E)$, every $j \in \mathbb{N}$ and $\mathbf{x} \in (V)_k$, if $|V| \geq 36$ then

$$\mathbb{P} [D_j(\mathbf{x}) > t \mid \text{Fill}^{\mathbf{x}}] \leq \exp \left\{ j - \frac{t}{\kappa_1 T_{\text{EX}}(4, f, G)(1/4)} \right\}.$$

Proof. Thanks to Lemma 7.3, the proofs of Lemmas 6.2 and 9.2 of [15] can be emulated to show that there exists a positive constant κ such that $\mathbb{E}[e^{D_j(\mathbf{x})/\kappa T} \mid \text{Fill}^{\mathbf{x}}] \leq e^j$ for all $j \in \mathbb{N}$. Thus by Markov’s inequality,

$$\begin{aligned} \mathbb{P} [D_j(\mathbf{x}) > t \mid \text{Fill}^{\mathbf{x}}] &= \mathbb{P} \left[e^{D_j(\mathbf{x})/\kappa T} > e^{t/\kappa T} \mid \text{Fill}^{\mathbf{x}} \right] \\ &\leq e^{-t/\kappa T} \mathbb{E}[e^{D_j(\mathbf{x})/\kappa T} \mid \text{Fill}^{\mathbf{x}}] \leq e^{j-t/\kappa T}. \end{aligned}$$

Writing $\kappa_1 = 20\kappa$ completes the proof. □

Combining Lemmas 7.2 and 7.4 completes the proof of part 3 of Lemma 3.1.

It therefore only remains to show that the chameleon process also satisfies part 2 of Lemma 3.1.

Let $\{\bar{\tau}_n\}_{n \in \mathbb{N}}$ denote the update times of the chameleon process $\{M_t\}_{t \geq 0}$; thus each $\bar{\tau}_n$ is either an incident time of the Poisson process Λ from Section 6.2, or a depinking time (of the form $2iT$ with $i \in \mathbb{N}$, as in Box 6.11). For each $j \in \mathbb{N}$, consider a process $\{M_t^j\}_{t \geq 0}$ which is identical to $\{M_t\}_{t \geq 0}$ for all $t < \bar{\tau}_j$ but evolves as the interchange process (i.e. with no further recolourings) for all $t \geq \bar{\tau}_j$. More formally, for all $t \geq \bar{\tau}_j$,

$$M_t^j = (I_{(\bar{\tau}_j, t]}(\mathbf{z}_{\bar{\tau}_j}), I_{(\bar{\tau}_j, t]}(R_{\bar{\tau}_j}), I_{(\bar{\tau}_j, t]}(P_{\bar{\tau}_j}), I_{(\bar{\tau}_j, t]}(W_{\bar{\tau}_j})),$$

where I is the map used in the modified graphical construction of the interchange process $\{\mathbf{x}_t^{\text{IP}}\}$ (see Section 6.2).

Notice that the almost-sure limit of $\{M_t^j\}_{t \geq 0}$ as $j \rightarrow \infty$ is the chameleon process $\{M_t\}_{t \geq 0}$. As a result, by the dominated convergence theorem, it suffices to prove that for each $j \in \mathbb{N}$ and $b \in V$,

$$\mathbb{P}[x_t^{\text{IP}} = b \mid \mathbf{z}_t^{\text{IP}}] = \mathbb{E}[\text{ink}_t^j(b) \mid \mathbf{z}_t^{\text{IP}}],$$

where $\text{ink}_t^j(b)$ is the amount of ink at vertex b in the process M_t^j . We prove this by induction on j . The case $j = 1$ is trivial since the particle initially at x is the only red particle (and there are no pink particles). For the inductive step we wish to show that almost surely

$$\mathbb{E}[\text{ink}_t^j(b) \mid \mathbf{z}_t^{\text{IP}}] = \mathbb{E}[\text{ink}_t^{j+1}(b) \mid \mathbf{z}_t^{\text{IP}}].$$

For $t < \bar{\tau}_j$, these are equal since the two processes evolve identically for such times. The update at time $\bar{\tau}_j$ of process $\{M_t^{j+1}\}$ is a chameleon step and could be of two types: also an update of the interchange process (i.e. $\bar{\tau}_j$ is an incident time of the Poisson process Λ), or not (i.e. it is a depinking time). Suppose we are in the first case. We condition on the common state of M^j and M^{j+1} at time $\bar{\tau}_{j-1}$. We want to show that almost surely

$$\mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^j(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^j] \right] = \mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^{j+1}(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}] \right]. \tag{7.1}$$

By the strong Markov property at time $\bar{\tau}_{j-1}$ we can construct a chameleon process $\{\tilde{M}_t^j\}$ (with associated interchange process $\tilde{\mathbf{x}}^{\text{IP}}$) which is identical to $\{M_t\}$ for all $t < \bar{\tau}_j$, but for all $t \geq \bar{\tau}_j$ evolves as an interchange process (i.e. with no further recolourings) and uses permutation choices:

- σ_n if $t = \tau_n$ is in a constant-colour phase,
- $\tilde{\beta}_{A(\sigma_n)}(\sigma_n)$ if $t = \tau_n$ is in a colour-changing phase.

We claim that

$$\frac{1}{2} \mathbb{E}[\text{ink}_{\bar{\tau}_j}^j(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^j] + \frac{1}{2} \mathbb{E}[\tilde{\text{ink}}_{\bar{\tau}_j}^j(b) \mid \tilde{\mathbf{z}}_{\bar{\tau}_j}^{\text{IP}}, \tilde{M}_{\bar{\tau}_{j-1}}^j] = \mathbb{E}[\text{ink}_{\bar{\tau}_j}^{j+1}(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}],$$

for all $b \in V$, almost surely (where $\tilde{\text{ink}}$ is the ink process under \tilde{M}^j). If $\bar{\tau}_j$ is in a constant-colour phase, then the statement is immediate (since all three processes update in exactly the same way). If $\bar{\tau}_j$ is in a colour-changing phase and the particle which is at b at time $\bar{\tau}_j$ has just been pinkened in the chameleon process then $\text{ink}_{\bar{\tau}_j}^{j+1}(b) = 1/2$ and by Lemma 6.12, $\{\text{ink}_{\bar{\tau}_j}^j(b), \tilde{\text{ink}}_{\bar{\tau}_j}^j(b)\} = \{0, 1\}$, and so the statement is true. Finally, if $\bar{\tau}_j$ is in a colour-changing phase but the particle at b at time $\bar{\tau}_j$ has not just been pinkened, then the three expectations are all equal since σ_j and $\tilde{\beta}_{A(\sigma_j)}(\sigma_j)$ have the same distribution, by Corollary 6.9 and Lemma 6.10 (and black particles move identically under each by Corollary 6.8). We thus have

$$\begin{aligned} \mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^j(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^j] \right] &= \frac{1}{2} \mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^j(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^j] \right] + \frac{1}{2} \mathbb{E} \left[\mathbb{E}[\tilde{\text{ink}}_{\bar{\tau}_j}^j(b) \mid \tilde{\mathbf{z}}_{\bar{\tau}_j}^{\text{IP}}, \tilde{M}_{\bar{\tau}_{j-1}}^j] \right] \\ &= \mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^{j+1}(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}] \right]. \end{aligned}$$

We are left to deal with the second case, when $\bar{\tau}_j$ is not an update of the interchange process. In this case there must be a depinking at time $\bar{\tau}_j$. We wish to show (7.1) holds, so again use the strong Markov property at time $\bar{\tau}_{j-1}$ to obtain

$$\begin{aligned} \mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^{j+1}(b) \mid \mathbf{z}_{\bar{\tau}_j}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}] \right] &= \mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^{j+1}(b) \mid \mathbf{z}_{\bar{\tau}_{j-1}}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}] \right] \\ &= \mathbb{E} \left[\mathbb{E}[\text{ink}_{\bar{\tau}_j}^j(b) \mid \mathbf{z}_{\bar{\tau}_{j-1}}^{\text{IP}}, M_{\bar{\tau}_{j-1}}^{j+1}] \right], \end{aligned}$$

where the second equality follows from the fact that an independent Bernoulli(1/2) random variable is used to determine the outcome of a depinking. This completes the induction, and with it the proof of part 2 of Lemma 3.1.

7.1 Bounding the rate of pinkening

In order to prove Lemma 7.3, we need to show that during a colour-changing phase (started with more white particles than red particles) the number of pink particles we create is in expectation at least a constant times the number of red particles at the start of that phase. We prove this result in this section.

Suppose we wish to lower-bound the number of white particles that are pinkened (which we shall refer to as the *number of pinkenings*) during the first colour-changing phase $[T, 2T]$. Since we start with 1 red particle, there will be more white particles at time T than red. We will wish to apply the following analysis for a general colour-changing phase (and not just the first) but the calculations will carry through since we are assuming the number of white particles is at least the number of red.

We make a change to the chameleon process in this section in order to ease our analysis – we remove the condition that we only pinken if we have fewer pink particles than either red or white particles and replace the set \bar{L}_t in the formal description (Box 6.11) with the potentially larger set L_t . Although this means we can end up with more pinkening events, this will only happen if a certain number of pinkening events have already happened (since pink particles are only created at times of pinkening events), and in that case we will be happy regardless. We shall refer to this new process as the *modified chameleon process*.

Let $a \in V$. We define t_a to be the smallest integer n such that $T < \tau_n \leq 2T$ and $a \in e_n$. If no such n exists we set $t_a = \infty$. Also, we set $\phi_a = \tau_{t_a}$ with notation $\tau_\infty = \infty$; hence ϕ_a is the first time (after time T) that vertex a is in a ringing edge of the underlying Poisson process. We define a third variable, F_a , set to be equal to $*$ in the case $\phi_a = \infty$. If, on the other hand, $\phi_a < \infty$, there are five possible cases. Let c_{t_a} denote the cycle of σ_{t_a} containing vertex a and $|c_{t_a}|$ denote the number of elements in c_{t_a} . For ease of notation we write d' for $\lfloor \frac{|c_{t_a}|}{4} \rfloor$ and m for the smallest element in c_{t_a} .

1. If $|c_{t_a}| = 2$ and $|e_{t_a}| = 2$, then set $F_a = I_{[T, \phi_a]}^{-1}(c_{t_a}(a))$.

2. If $|c_{t_a}| = 2$ and $|e_{t_a}| \geq 3$, then denote by

$$(a_1 \ a_2) \prec \dots \prec (a_{l-1} \ a_l)$$

the ordered transpositions in the canonical cyclic decomposition of ρ_{t_a} . If there exists $j \in \{1, \dots, \lfloor l/4 \rfloor\}$ with $a \in \{a_{4j-3}, a_{4j-2}, a_{4j-1}, a_{4j}\}$ and $a_{4j-1} < a_{4j-2}$, then:

(a). if $a = a_{4j-3}$ set $F_a = I_{[T, \phi_a]}^{-1}(a_{4j-1}, a_{4j-2}, a_{4j})$;

(b). if $a = a_{4j-2}$ set $F_a = I_{[T, \phi_a]}^{-1}(a_{4j}, a_{4j-3}, a_{4j-1})$;

(c). if $a = a_{4j-1}$ set $F_a = I_{[T, \phi_a]}^{-1}(a_{4j-3}, a_{4j-2}, a_{4j})$;

(d). if $a = a_{4j}$ set $F_a = I_{[T, \phi_a]}^{-1}(a_{4j-2}, a_{4j-3}, a_{4j-1})$.

3. If $|c_{t_a}| = 3$, then set $F_a = I_{[T, \phi_a]}^{-1}(c_{t_a}(a), c_{t_a}^2(a))$.

4. If $|c_{t_a}| \geq 4$, and there exists $j \in \{1, \dots, d'\}$ with $a \in c_{t_a}^{H_j}(m)$, then:

(i). if $a = c_{t_a}^{2j-2}(m)$, set

$$F_a = I_{[T, \phi_a]}^{-1}(c_{t_a}^{2d'+2j-2}(m), c_{t_a}^{2j-1}(m), c_{t_a}^{2d'+2j-1}(m));$$

(ii). if $a = c_{t_a}^{2d'+2j-2}(m)$, set

$$F_a = I_{[T, \phi_a]}^{-1}(c_{t_a}^{2j-2}(m), c_{t_a}^{2j-1}(m), c_{t_a}^{2d'+2j-1}(m));$$

(iii). if $a = c_{t_a}^{2j-1}(m)$, set

$$F_a = I_{[T, \phi_a]}^{-1}(c_{t_a}^{2d'+2j-1}(m), c_{t_a}^{2j-2}(m), c_{t_a}^{2d'+2j-2}(m));$$

(iv). if $a = c_{t_a}^{2d'+2j-1}(m)$, set

$$F_a = I_{[T, \phi_a]}^{-1}(c_{t_a}^{2j-1}(m), c_{t_a}^{2j-2}(m), c_{t_a}^{2d'+2j-2}(m)).$$

5. In all other cases, set $F_a = *$.

Remark 7.5. From this construction it is easy to see that (for any $b, c, d \in V$)

1. In case 1 above, we have

$$\{F_a = b, \phi_a = \phi_b\} = \{F_b = a, \phi_a = \phi_b\}.$$

2. In case 3 above, we have

$$\begin{aligned} \{F_a = (b, c), \phi_a = \phi_b = \phi_c\} &= \{F_b = (c, a), \phi_a = \phi_b = \phi_c\} \\ &= \{F_c = (a, b), \phi_a = \phi_b = \phi_c\}. \end{aligned}$$

3. In all other cases, we have

$$\begin{aligned} \{F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d\} &= \{F_b = (a, c, d), \phi_a = \phi_b = \phi_c = \phi_d\} \\ &= \{F_c = (d, a, b), \phi_a = \phi_b = \phi_c = \phi_d\} = \{F_d = (c, a, b), \phi_a = \phi_b = \phi_c = \phi_d\}. \end{aligned}$$

A possible evolution of the chameleon process during the first two phases is shown in Figure 9.

We now present a method to count the number of pinkenings during a colour-changing phase of the modified chameleon process. For ease of notation we shall write I for the map $I_{[0, T]}$. The proofs of the first three results below are fairly simple extensions of equivalent results in [15] and can be found in Appendix C.

Lemma 7.6. Consider a modified chameleon process with its starting configuration (z, R, P, W) satisfying $|P| < |R| \leq |W|$. Then the number of pinkenings during $(T, 2T)$ is at least the number of $b \in I(W)$ such that one of the following holds:

- $F_a = b$ for some $a \in I(R)$ with $\phi_a = \phi_b$,
- $F_a = (b, c)$ for some $a \in I(R)$ and $c \in I(W)$ with $\phi_a = \phi_b = \phi_c$, and $\theta_{t_a} = 1$,
- $F_a = (b, c, d)$ for some $a \in I(R)$ and $c, d \in I(W)$ with $\phi_a = \phi_b = \phi_c = \phi_d$, and $\theta_{t_a} = 1$.

In bounding the expected number of pinkenings during a colour-changing phase, it turns out to be useful to have a lower bound on the probability that $F_a \neq *$ given $\phi_a \neq \infty$. This is because even if vertex a is in a ringing edge during time interval $[T, 2T]$, in order for the particle initially at a to be pinkened in the modified chameleon process at this time, it is necessary (but not sufficient) that $F_a \neq *$. The proof of this lemma makes use of part 3 of Assumption 1.1.

Lemma 7.7. For every $a \in V$, $\mathbb{P}[F_a \neq * | \phi_a \neq \infty] \geq 4/15$.

Proposition 7.8. Consider a modified chameleon process with initial configuration (z, R, P, W) . Then for any vertices a, b, c, d ,

$$(i). \mathbb{P}[|\{a, b\} \cap I(R)| = 1, |\{a, b\} \cap I(W)| = 1] \geq (1 - 2^{-9})^2 |R| \frac{|W|}{\binom{|V|}{2}},$$




Figure 9: In this example we suppose that $V = \{1, 2, 3, 4, 5, 6\}$, $E = \{\{1, 2, 3, 4\}, \{1, 2, 5, 6\}, \{3, 4, 5, 6\}\}$, and permutations chosen are uniform 4-cycles. Particles are labelled according to their initial location. We see that $I_T(R) = I_T(\{1\}) = \{4\}$, $t_4 = 3$ (the first time vertex 4 is in a ringing edge after time T is τ_3), $\tau_3 = \phi_1 = \phi_2 = \phi_3 = \phi_4$ (the first time vertices 1, 2, 3 and 4 are in a ringing edge after time T is τ_3), and $F_4 = (2, 1, 3)$. Particles at vertices 2 and 4 (at time T) are coloured pink at time τ_3 .

$$(ii). \mathbb{P} [|\{a, b, c\} \cap I(R)| = 1, |\{a, b, c\} \cap I(W)| = 2] \geq (1 - 2^{-9})^2 |R| \frac{\binom{|W|}{2}}{\binom{|V|}{3}},$$

$$(iii). \mathbb{P} [|\{a, b, c, d\} \cap I(R)| = 1, |\{a, b, c, d\} \cap I(W)| = 3] \geq (1 - 2^{-9})^2 |R| \frac{\binom{|W|}{3}}{\binom{|V|}{4}}.$$

We now present the main result of this section – a version of Lemma 7.3 but proved for the *modified* chameleon process. As explained earlier in this section, this implies the corresponding result for our original chameleon process.

Lemma 7.9. *Suppose $|V| \geq 36$ and consider a modified chameleon process with initial configuration (z, R, P, W) satisfying $|P| < |R| \leq |W|$. Then*

$$\mathbb{E}[|R_{2T-}|] \leq (1 - 10^{-6})|R|.$$

Proof. Write $N(b)$ for the set of vertices that share at least one edge of the hypergraph with b . By Lemma 7.6, we have

$$\begin{aligned}
 & |R_{2T-}| \\
 & \leq |R| - \sum_{b \in I(W)} \mathbf{1}_{\{\cup_{a \in N(b)} \{F_a = b, \phi_a = \phi_b, a \in I(R)\}\}} \\
 & \quad - \sum_{b \in I(W)} \mathbf{1}_{\{\cup_{a, c \in N(b)} \{F_a = (b, c), \phi_a = \phi_b = \phi_c, a \in I(R), c \in I(W), \theta_{t_a} = 1\}\}} \\
 & \quad - \sum_{b \in I(W)} \mathbf{1}_{\{\cup_{a, c, d \in N(b)} \{F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d, a \in I(R), c, d \in I(W), \theta_{t_a} = 1\}\}} \\
 & = |R| - \sum_{b \in I(W)} \sum_{a \in N(b)} \left\{ \mathbf{1}_{\{F_a = b, \phi_a = \phi_b, a \in I(R)\}} \right. \\
 & \quad + \sum_{c \in N(b)} \left\{ \mathbf{1}_{\{F_a = (b, c), \phi_a = \phi_b = \phi_c, a \in I(R), c \in I(W), \theta_{t_a} = 1\}} \right. \\
 & \quad \left. \left. + \sum_{d \in N(b)} \mathbf{1}_{\{F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d, a \in I(R), c, d \in I(W), \theta_{t_a} = 1\}} \right\} \right\}.
 \end{aligned}$$

Note now that the event $\{F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d\}$ is determined entirely by the process after time T , and in particular is independent of the process between times 0 and T , and hence of the map $I = I_{[0, T]}$. This is also true for the event $\{F_a = (b, c), \phi_a = \phi_b = \phi_c\}$ and the event $\{F_a = b, \phi_a = \phi_b\}$. Recalling Remark 7.5 we see that the expectation of the above is equal to

$$\begin{aligned}
 & |R| - \sum_{b \in V} \sum_{a \in N(b)} \left\{ \frac{1}{2} \mathbb{P} [F_a = b, \phi_a = \phi_b] \mathbb{P} [|\{a, b\} \cap I(R)| = 1, |\{a, b\} \cap I(W)| = 1] \right. \\
 & \quad + \mathbb{P} [\theta_{t_a} = 1] \sum_{c \in N(b)} \left\{ \frac{1}{3} \mathbb{P} [F_a = (b, c), \phi_a = \phi_b = \phi_c] \right. \\
 & \quad \quad \cdot \mathbb{P} [|\{a, b, c\} \cap I(R)| = 1, |\{a, b, c\} \cap I(W)| = 2] \\
 & \quad \left. + \sum_{d \in N(b)} \frac{1}{4} \mathbb{P} [F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d] \right. \\
 & \quad \quad \left. \left. \mathbb{P} [|\{a, b, c, d\} \cap I(R)| = 1, |\{a, b, c, d\} \cap I(W)| = 3] \right\} \right\}.
 \end{aligned}$$

Using Proposition 7.8 and $\mathbb{P} [\theta_{t_a} = 1] = 1/2$ we obtain the bound

$$\begin{aligned}
 \mathbb{E} [|R_{2T-}|] - |R| & \leq - (1 - 2^{-9})^2 \sum_{b \in V} \sum_{a \in N(b)} \left\{ \frac{1}{2} \frac{|R| |W|}{\binom{|V|}{2}} \mathbb{P} [F_a = b, \phi_a = \phi_b] \right. \tag{7.2} \\
 & \quad + \sum_{c \in N(b)} \left\{ \frac{1}{6} \frac{|R| \binom{|W|}{2}}{\binom{|V|}{3}} \mathbb{P} [F_a = (b, c), \phi_a = \phi_b = \phi_c] \right. \\
 & \quad \left. \left. + \sum_{d \in N(b)} \frac{1}{8} \frac{|R| \binom{|W|}{3}}{\binom{|V|}{4}} \mathbb{P} [F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d] \right\} \right\}.
 \end{aligned}$$

Consider the final probability in the above equation. We can write it as

$$\begin{aligned}
 & \mathbb{P} [F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d] \\
 & = \sum_{\substack{e \in E: \\ a, b, c, d \in e}} \mathbb{P} [F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d, e_{t_a} = e]
 \end{aligned}$$

$$\begin{aligned}
 &= \sum_{\substack{e \in E: \\ a,b,c,d \in e}} \mathbb{P} [F_a = (b, c, d) \mid |F_a| = 3, e_{t_a} = e] \mathbb{P} [|F_a| = 3, \phi_a = \phi_b = \phi_c = \phi_d, e_{t_a} = e] \\
 &= \sum_{\substack{e \in E: \\ a,b,c,d \in e}} \mathbb{P} [F_a = (b, c, d) \mid |F_a| = 3, e_{t_a} = e] \mathbb{P} [|F_a| = 3] \mathbb{P} [\phi_a = \phi_b = \phi_c = \phi_d, e_{t_a} = e],
 \end{aligned}$$

where we have made use of the fact that the choice of the permutations (which determines $|F_a|$) is independent of the choice of the edges that ring. We next make use of the regularity of the hypergraph (and that all edges ring at the same rate) to obtain that $\mathbb{P} [\phi_a = \phi_b = \phi_c = \phi_d, e_{t_a} = e] \geq 1/(4D)$ where D is the degree of each vertex. Summing the above over a, b, c, d gives

$$\sum_{b \in V} \sum_{a,c,d \in N(b)} \mathbb{P} [F_a = (b, c, d), \phi_a = \phi_b = \phi_c = \phi_d] \geq \frac{1}{4D} \sum_{e \in E} \sum_{a \in e} \mathbb{P} [|F_a| = 3]. \tag{7.3}$$

Similarly

$$\sum_{b \in V} \sum_{a,c \in N(b)} \mathbb{P} [F_a = (b, c), \phi_a = \phi_b = \phi_c] \geq \frac{1}{3D} \sum_{e \in E} \sum_{a \in e} \mathbb{P} [|F_a| = 2], \tag{7.4}$$

and

$$\sum_{b \in V} \sum_{a \in N(b)} \mathbb{P} [F_a = b, \phi_a = \phi_b] \geq \frac{1}{2D} \sum_{e \in E} \sum_{a \in e} \mathbb{P} [|F_a| = 1]. \tag{7.5}$$

Combining (7.2), (7.3), (7.4) and (7.5) gives

$$\mathbb{E} [|R_{2T-}|] - |R| \leq -\frac{1}{4D} (1 - 2^{-9})^2 \frac{1}{8} \frac{|R| \binom{|W|}{3}}{\binom{|V|}{4}} \sum_{a \in V} \sum_{e: e \ni a} \mathbb{P} [F_a \neq *].$$

Using Lemma 7.7,

$$\mathbb{P} [F_a \neq *] = \mathbb{P} [\phi_a \neq \infty] \mathbb{P} [F_a \neq * \mid \phi_a \neq \infty] \geq \frac{4}{15} \mathbb{P} [\phi_a \neq \infty]. \tag{7.6}$$

Also,

$$\mathbb{P} [\phi_a = \infty] \leq \mathbb{P} [I_{(T,2T)}(a) = a] = \mathbb{P} [a_T^{\text{RW}} = a],$$

where $\{a_t^{\text{RW}}\}$ is a realisation of $\text{RW}(1, f, G)$ started from a . Since

$$T = 20T_{\text{IP}(4,f,G)}(1/4) \geq 20T_{\text{RW}(f,G)}(1/4) \geq T_{\text{RW}(f,G)}(2^{-20})$$

(by Proposition 2.2) we have

$$\mathbb{P} [\phi_a = \infty] \leq \mathbb{P} [a_T^{\text{RW}} = a] \leq \frac{1}{|V|} + 2^{-20}.$$

From (7.6), we deduce that

$$\mathbb{P} [F_a \neq *] \geq \frac{4}{15} (1 - 2^{-20} - 1/|V|). \tag{7.7}$$

Finally, the assumptions in Lemma 7.3 on the sizes of the sets P, R and W imply that

$$3|W| \geq |W| + |R| + |P| = |V| - k + 1 \geq |V|/2,$$

and since $|V| \geq 36$ we arrive at our stated result:

$$\begin{aligned}
 \mathbb{E} [|R_{2T-}|] - |R| &\leq -\frac{1}{4D} (1 - 2^{-9})^2 \frac{|R|}{864|V|} \sum_{a \in V} \sum_{e: e \ni a} \frac{4}{15} (1 - 2^{-20} - 1/|V|) \\
 &= -\frac{1}{4D} (1 - 2^{-9})^2 \frac{|R|}{864|V|} |V|D \frac{4}{15} (1 - 2^{-20} - 1/|V|) \\
 &\leq -10^{-6}|R|. \quad \square
 \end{aligned}$$

A Case by case analysis from Proof of Lemma 4.4




Figure 10: The bijection between $\sigma_{a,b}$ and $\tilde{\sigma}_{a,b}$. We see that there are certain failure permutations, for which no match occurs, where a match refers to the event that particles a^* and b^* are moved to the same location after applying permutations $\sigma_{a,b}$ and $\tilde{\sigma}_{a,b}$, respectively. For a fixed cycle structure we see that the probability of ‘no match’ is at most $1/2$, which is achieved when the edge-size is 3 and a cycle of size 3 is chosen.

B Technical proofs for Section 4

Here we include some of the more technical proofs required to compare the mixing time of $EX(4, f, G)$ with that of $EX(2, f, G)$.

Proof of Lemma 4.5. Since the hypergraph is non-easy there exists $\mathbf{x} \in V^2$ such that

$$\mathbb{P} [M^{\text{RW}}(\mathbf{x}) > 10^{10}T] > 1/1000.$$

We have

$$\begin{aligned} & \mathbb{P} [M^{\text{RW}}(\mathbf{x}) > 10^{10}T] \\ &= \mathbb{E} [\mathbb{E} [\mathbf{1}_{\{M^{\text{RW}}(\mathbf{x}) > 10^{10}T\}} | \mathbf{x}_{(10^{10}-40)T}^{\text{RW}}]] \\ &= \mathbb{E} \left[\mathbb{E} \left[\mathbf{1}_{\{M^{\text{RW}}(\mathbf{x}) > (10^{10}-40)T\}} | \mathbf{x}_{(10^{10}-40)T}^{\text{RW}} \right] \mathbb{E} \left[\mathbf{1}_{\{M^{\text{RW}}(\mathbf{x}_{(10^{10}-40)T}^{\text{RW}}) > 40T\}} | \mathbf{x}_{(10^{10}-40)T}^{\text{RW}} \right] \right] \\ &\leq \mathbb{E} \left[\mathbb{E} \left[\mathbf{1}_{\{M^{\text{RW}}(\mathbf{x}) > (10^{10}-40)T\}} | \mathbf{x}_{(10^{10}-40)T}^{\text{RW}} \right] \sup_{\mathbf{y} \in V^2} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 40T] \right] \\ &= \sup_{\mathbf{y} \in V^2} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 40T] \mathbb{P} [M^{\text{RW}}(\mathbf{x}) > (10^{10} - 40)T] \\ &\leq \left(\sup_{\mathbf{y} \in V^2} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 40T] \right)^{\frac{10^{10}}{40}}. \end{aligned}$$

Hence there exists $\mathbf{y} \in V^2$ such that

$$\mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 40T] \geq \left(\frac{1}{1000} \right)^{40(10^{-10})} > 1 - 10^{-7}.$$

Now,

$$\begin{aligned} \mathbb{P} [M^{\text{RW}}(\mathbf{y}) > 40T] &\leq \mathbb{P} [M^{\text{RW}}(\mathbf{y}_{20T}^{\text{RW}}) > 20T] \\ &= \sum_{\mathbf{v} \in V^2} \mathbb{P} [\mathbf{y}_{20T}^{\text{RW}} = \mathbf{v}] \mathbb{P} [M^{\text{RW}}(\mathbf{v}) > 20T]. \end{aligned}$$

However, by Definition (2.3) of total-variation, Proposition 2.2, and the fact that (by the contraction principle) $T_{\text{RW}(2,f,G)}(1/4) \leq T_{\text{EX}(2,f,G)}(1/4)$,

$$\begin{aligned} \sum_{\mathbf{v} \in V^2} \mathbb{P} [\mathbf{y}_{20T}^{\text{RW}} = \mathbf{v}] \mathbb{P} [M^{\text{RW}}(\mathbf{v}) > 20T] &- \sum_{\mathbf{v} \in V^2} \frac{\mathbb{P} [M^{\text{RW}}(\mathbf{v}) > 20T]}{|V|^2} \\ &\leq \|\mathcal{L}[\mathbf{y}_{20T}^{\text{RW}}] - \text{Unif}(V^2)\|_{\text{TV}} \leq 2^{-20}. \end{aligned}$$

Hence

$$\begin{aligned} \sum_{\mathbf{v} \in V^2} \frac{\mathbb{P} [M^{\text{RW}}(\mathbf{v}) > 20T]}{|V|^2} &\geq \sum_{\mathbf{v} \in V^2} \mathbb{P} [\mathbf{y}_{20T}^{\text{RW}} = \mathbf{v}] \mathbb{P} [M^{\text{RW}}(\mathbf{v}) > 20T] - 2^{-20} \\ &\geq 1 - 10^{-7} - 2^{-20} \\ &\geq 1 - \frac{1}{1000}. \end{aligned} \quad \square$$

Proof of Lemma 4.6. We begin by conditioning on the value of $\mathbf{O}(\mathbf{x})_{20T}^{\text{EX}}$.

$$\begin{aligned} & \mathbb{P} [\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x})_{20T}^{\text{EX}}) \leq 20T] \\ &= \sum_{\{a_1, a_2, a_3, a_4\} \in \binom{V}{4}} \mathbb{P} [\bar{M}^{\text{RW}}(\{a_1, a_2, a_3, a_4\}) \leq 20T] \mathbb{P} [\mathbf{O}(\mathbf{x})_{20T}^{\text{EX}} = \{a_1, a_2, a_3, a_4\}]. \end{aligned}$$

For each $\mathbf{a} \in \binom{V}{4}$, we have

$$\mathbb{P} [\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{a})) \leq 20T] \leq \sum_{i=2}^4 \sum_{j=1}^{i-1} \mathbb{P} [M^{\text{RW}}((\mathbf{a}(i), \mathbf{a}(j))) \leq 20T],$$

and so

$$\begin{aligned}
 & \mathbb{P} [\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x})_{20T}^{\text{EX}}) \leq 20T] \\
 & \leq \sum_{\{u_1, u_2\} \in \binom{V}{2}} \mathbb{P} [M^{\text{RW}}((u_1, u_2)) \leq 20T] \mathbb{P} [\mathbf{O}(\mathbf{x})_{20T}^{\text{EX}} \supset \{u_1, u_2\}] \\
 & \leq \sum_{i=2}^4 \sum_{j=1}^{i-1} \sum_{\{u_1, u_2\} \in \binom{V}{2}} \mathbb{P} [M^{\text{RW}}((u_1, u_2)) \leq 20T] \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{u_1, u_2\}]. \tag{B.1}
 \end{aligned}$$

Now, for each $1 \leq j < i \leq 4$ (and motivated by [15]) set

$$\text{Good}_{i,j} := \left\{ \{a, b\} \in \binom{V}{2} : \left| \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{a, b\}] - \frac{1}{\binom{|V|}{2}} \right| \leq \frac{\varepsilon}{\binom{|V|}{2}} \right\}.$$

We decompose the sum over \mathbf{u} above into $\mathbf{u} \in \text{Good}_{i,j}$ and $\mathbf{u} \in \text{Bad}_{i,j}$, where

$$\text{Bad}_{i,j} = \binom{V}{2} \setminus \text{Good}_{i,j}.$$

For the Good terms, we have

$$\begin{aligned}
 & \sum_{i=2}^4 \sum_{j=1}^{i-1} \sum_{\{u_1, u_2\} \in \text{Good}_{i,j}} \mathbb{P} [M^{\text{RW}}((u_1, u_2)) \leq 20T] \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{u_1, u_2\}] \\
 & \leq \sum_{i=2}^4 \sum_{j=1}^{i-1} \sum_{\{u_1, u_2\} \in \text{Good}_{i,j}} \mathbb{P} [M^{\text{RW}}((u_1, u_2)) \leq 20T] \frac{(1 + \varepsilon)}{\binom{|V|}{2}} \\
 & \leq 25(1 + \varepsilon) \sum_{\mathbf{u} \in V^2} \frac{\mathbb{P} [M^{\text{RW}}(\mathbf{u}) \leq 20T]}{|V|^2}. \tag{B.2}
 \end{aligned}$$

For the Bad terms, we have

$$\begin{aligned}
 & \sum_{i=2}^4 \sum_{j=1}^{i-1} \sum_{\{u_1, u_2\} \in \text{Bad}_{i,j}} \mathbb{P} [M^{\text{RW}}((u_1, u_2)) \leq 20T] \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{u_1, u_2\}] \\
 & \leq \sum_{i=2}^4 \sum_{j=1}^{i-1} \sum_{\{u_1, u_2\} \in \text{Bad}_{i,j}} \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{u_1, u_2\}] \\
 & = \sum_{i=2}^4 \sum_{j=1}^{i-1} \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} \in \text{Bad}_{i,j}]. \tag{B.3}
 \end{aligned}$$

Note that for each $1 \leq j < i \leq 4$,

$$\begin{aligned}
 \|\mathcal{L}(\mathbf{x}_{20T}^{\text{EX}}) - \text{Uniform}\|_{\text{TV}} &= \frac{1}{2} \sum_{\{u_1, u_2\} \in \binom{V}{2}} \left| \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{u_1, u_2\}] - \frac{1}{\binom{|V|}{2}} \right| \\
 &> \frac{1}{2} \frac{\varepsilon}{\binom{|V|}{2}} |\text{Bad}_{i,j}|,
 \end{aligned}$$

since every $\{u_1, u_2\} \in \text{Bad}_{i,j}$ contributes at least $\varepsilon/\binom{|V|}{2}$ to the sum. However, the left-hand side in the above equation is at most 2^{-20} by the choice of T . We deduce that

$$|\text{Bad}_{i,j}| \leq \varepsilon^{-1} 2^{-19} \binom{|V|}{2},$$

and thus

$$|\text{Good}_{i,j}| \geq (1 - \varepsilon^{-1}2^{-19}) \binom{|V|}{2}.$$

However, for each $\{u_1, u_2\} \in \text{Good}_{i,j}$ we know that

$$\mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{u_1, u_2\}] \geq \frac{1 - \varepsilon}{\binom{|V|}{2}}.$$

Therefore,

$$\mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} \in \text{Good}_{i,j}] \geq \frac{1 - \varepsilon}{\binom{|V|}{2}} |\text{Good}_{i,j}| \geq 1 - \varepsilon - \varepsilon^{-1}2^{-19}.$$

We deduce that

$$\mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} \in \text{Bad}_{i,j}] \leq \varepsilon + \varepsilon^{-1}2^{-19}.$$

Plugging this into (B.3) gives

$$\begin{aligned} & \sum_{i=2}^4 \sum_{j=1}^{i-1} \sum_{\{u_1, u_2\} \in \text{Bad}_{i,j}} \mathbb{P} [M^{\text{RW}}((u_1, u_2)) \leq 20T] \mathbb{P} [\{\mathbf{x}(i), \mathbf{x}(j)\}_{20T}^{\text{EX}} = \{u_1, u_2\}] \\ & \leq 12(\varepsilon + \varepsilon^{-1}2^{-19}). \end{aligned}$$

Combining this with (B.2) and (B.1) gives

$$\mathbb{P} [\bar{M}^{\text{RW}}(\mathbf{x}_{20T}^{\text{IP}}) \leq 20T] \leq 12(\varepsilon + \varepsilon^{-1}2^{-19}) + 25(1 + \varepsilon) \sum_{\mathbf{u} \in V^2} \frac{\mathbb{P} [M^{\text{RW}}(\mathbf{u}) \leq 20T]}{|V|^2}. \quad \square$$

Proof of Proposition 4.7. This proof is similar to the proof of Proposition 4.6 in [15]. By the graphical construction of Section 2.2, $\mathbf{O}(\mathbf{x})_s^{\text{EX}}$ and $\mathbf{O}(\mathbf{x}_s^{\text{IP}})$ have the same distribution. Thus by the contraction principle it suffices to show that

$$\|\mathcal{L}[\mathbf{x}_s^{\text{RW}}] - \mathcal{L}[\mathbf{x}_s^{\text{IP}}]\|_{\text{TV}} \leq \mathbb{P} [\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x})) \leq s].$$

We present a coupling of $\{\mathbf{x}_t^{\text{IP}}\}_{t \geq 0}$ and $\{\mathbf{x}_t^{\text{RW}}\}_{t \geq 0}$ such that the two processes agree up to time $\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x}))$. The coupling has state-space $S := (V)_2 \times V^2$ which we split into two parts: $\Delta := \{(\mathbf{z}, \mathbf{z}) : \mathbf{z} \in (V)_2\}$ and $\Delta^{\mathbb{G}}$. Denote by $q(\cdot, \cdot)$ the transition rates. We construct the coupling as follows:

1. if $(\mathbf{x}, \mathbf{y}) \in \Delta^{\mathbb{G}}$, the transition rate to any other state in S is the same as that of independent realisations of $\text{IP}(4, f, G)$ and $\text{RW}(4, f, G)$.
2. if $(\mathbf{x}, \mathbf{x}) \in \Delta$,
 - (a) for $e \in E$ with $|e \cap \{x(1), x(2), x(3), x(4)\}| = 1$ and for each $\sigma_e \in \mathcal{S}_e$,

$$q((\mathbf{x}, \mathbf{x}), (\sigma_e(\mathbf{x}), \sigma_e(\mathbf{x}))) = f_e(\sigma_e).$$

- (b) for $e \in E$ with $|e \cap \{\mathbf{x}(1), \mathbf{x}(2), \mathbf{x}(3), \mathbf{x}(4)\}| > 1$ and for each $\sigma_e \in \mathcal{S}_e$,

$$q((\mathbf{x}, \mathbf{x}), (\sigma_e(\mathbf{x}), (\sigma_e(\mathbf{x}(1)), \mathbf{x}(2), \mathbf{x}(3), \mathbf{x}(4)))) = f_e(\sigma_e),$$

$$q((\mathbf{x}, \mathbf{x}), (\mathbf{x}, (\mathbf{x}(1), \sigma_e(\mathbf{x}(2)), \mathbf{x}(3), \mathbf{x}(4)))) = f_e(\sigma_e),$$

$$q((\mathbf{x}, \mathbf{x}), (\mathbf{x}, (\mathbf{x}(1), \mathbf{x}(2), \sigma_e(\mathbf{x}(3)), \mathbf{x}(4)))) = f_e(\sigma_e),$$

$$q((\mathbf{x}, \mathbf{x}), (\mathbf{x}, (\mathbf{x}(1), \mathbf{x}(2), \mathbf{x}(3), \sigma_e(\mathbf{x}(4)))) = f_e(\sigma_e).$$

3. all other transitions have rate 0.

By inspection of the marginals it is clear that this indeed gives a coupling of the two processes. Furthermore, if we start the coupling from a state $\mathbf{x} \in \Delta$, the two processes can only differ after a transition has occurred according to rule 2.(b); but the first time this happens is precisely $\bar{M}^{\text{RW}}(\mathbf{O}(\mathbf{x}))$. \square

C Technical proofs for Section 7

Here we include proofs of some of the results used in Section 7.1.

Proof of Lemma 7.6. We shall show that in each situation the particle at vertex b at time T is pinkened during $(T, 2T)$.

In the first situation with $F_a = b$ for some $a \in I(R)$ with $\phi_a = \phi_b$, we deduce that $|e_{t_a}| = 2$ and $I_{[T, \phi_a)}(b) = \sigma_{t_a}(a)$. Since $\phi_a = \phi_b$, we have

$$\sigma_{t_a}(a) = I_{[T, \phi_a)}(b) = I_{[T, \phi_b)}(b) = b.$$

Since $a \in I(R)$, $b \in I(W)$ and $\phi_a = \phi_b$, we have that $a \in I_{[0, \phi_a)}(R)$ and $b \in I_{[0, \phi_a)}(W)$. This implies that the particle at b at time T (and also the particle at a at time T) is pinkened at time ϕ_a .

In the second situation with $F_a = (b, c)$ for some $a \in I(R)$, $c \in I(W)$ with $\phi_a = \phi_b = \phi_c$ and $\theta_{t_a} = 1$, we deduce that $|c_{t_a}| = 3$ and $I_{[T, \phi_a)}(b, c) = (c_{t_a}(a), c_{t_a}^2(a))$. Since $\phi_a = \phi_b = \phi_c$, we have

$$(c_{t_a}(a), c_{t_a}^2(a)) = (b, c).$$

Since $a \in I(R)$, $b, c \in I(W)$ and $\phi_a = \phi_b = \phi_c$, we have that $a \in I_{[0, \phi_a)}(R)$ and $b, c \in I_{[0, \phi_a)}(W)$ and hence it is immediate that there exists $1 \leq i \leq K$ satisfying $A_i(R_{\phi_a-}, W_{\phi_a-}, c_{\phi_a}) = \{0\}$. Since $\theta_{t_a} = 1$ we deduce that the particle at b at time T is pinkened at time ϕ_a .

In the third situation with $F_a = (b, c, d)$ for some $a \in I(R)$, $c, d \in I(W)$ with $\phi_a = \phi_b = \phi_c = \phi_d$ and $\theta_{t_a} = 1$ we have two cases. The first case is if $|c_{t_a}| = 2$. We denote by

$$(a_1 \ a_2) \prec \cdots \prec (a_{l-1} \ a_l)$$

the ordered transpositions in the cyclic decomposition of σ_{t_a} and denote by ρ_0 the composition of these transpositions. There are four sub-cases which are all similar, and we just prove the result for one of them. So suppose there exists $j \in \{1, \dots, \lfloor l/4 \rfloor\}$ with $a = a_{4j-3}$. Then we have $I_{[T, \phi_a)}(b, c, d) = (a_{4j-1}, a_{4j-2}, a_{4j})$. Since $\phi_a = \phi_b = \phi_c = \phi_d$, we have

$$(a_{4j-1}, a_{4j-2}, a_{4j}) = (b, c, d).$$

Since $a \in I(R)$, $b, c, d \in I(W)$ and $\phi_a = \phi_b = \phi_c = \phi_d$, we have that $a \in I_{[0, \phi_a)}(R)$ and $b, c, d \in I_{[0, \phi_a)}(W)$ and hence $j \in A_0(R_{\phi_a-}, W_{\phi_a-}, \rho_0)$ with $L_{\phi_a}^{0,j} = \{a, b\}$. Since $\theta_{t_a} = 1$ we deduce that the particle at b at time T is pinkened at time ϕ_a . The other three sub-cases follow similarly.

The second possibility when $F_a = (b, c, d)$ is that $|c_{t_a}| \geq 4$. Again there are four sub-cases which are all similar, and we just prove the result for one of them. So suppose there exists $j \in \{1, \dots, d'\}$ with $a = c_{t_a}^{2j-2}(m)$. Then we have $I_{[T, \phi_a)}(b, c, d) = (c_{t_a}^{2d'+2j-2}(m), c_{t_a}^{2j-1}(m), c_{t_a}^{2d'+2j-1}(m))$. Since $\phi_a = \phi_b = \phi_c = \phi_d$, we have

$$(c_{t_a}^{2d'+2j-2}(m), c_{t_a}^{2j-1}(m), c_{t_a}^{2d'+2j-1}(m)) = (b, c, d).$$

Since $a \in I(R)$, $b, c, d \in I(W)$ and $\phi_a = \phi_b = \phi_c = \phi_d$, we have that $a \in I_{[0, \phi_a)}(R)$ and $b, c, d \in I_{[0, \phi_a)}(W)$ and hence $j \in A_i(R_{\phi_a-}, W_{\phi_a-}, \sigma_{t_a})$ for some $1 \leq i \leq K$ with

$L_{\phi_a}^{i,j} = \{a, b\}$. Since $\theta_{t_a} = 1$ we deduce that the particle at b at time T is pinkened at time ϕ_a . \square

Proof of Lemma 7.7. For a fixed $a \in V$ we wish to upper bound $\mathbb{P}[F_a = * | \phi_a \neq \infty]$. Recall that we write c_{t_a} for the cycle of σ_{t_a} containing vertex a , d' for $\lfloor \frac{|c_{t_a}|}{4} \rfloor$, and m for the smallest element in c_{t_a} . On the event $\{\phi_a \neq *\}$, the event $\{F_a = *\}$ is equivalent to the event that at time ϕ_a one of the following occurs:

Event B_1 : $|c_{t_a}| = 1$,

Event B_2 : $|e_{t_a}| \geq 3$, $|c_{t_a}| = 2$ and $l = 2$

(i.e. there is only one transposition and it contains vertex a),

Event B_3 : $|e_{t_a}| \geq 3$, $|c_{t_a}| = 2$, $l \neq 2$, $a \in \{a_{l-1}, a_l\}$ and $l \notin 4\mathbb{N}$,

Event B_4 : $|e_{t_a}| \geq 3$, $|c_{t_a}| = 2$, there exists $j \in \{1, \dots, \lfloor l/4 \rfloor\}$ with

$$a \in \{a_{4j-3}, a_{4j-2}, a_{4j-1}, a_{4j}\} \text{ and } a_{4j-1} > a_{4j-2},$$

Event B_5 : $|c_{t_a}| \geq 4$ but there does not exist a $j \in \{1, \dots, d'\}$ and

$$r \in \{2j - 2, 2d'(c_{t_a}) + 2j - 2, 2j - 1, 2d'(c_{t_a}) + 2j - 1\} \text{ with } a = c_{t_a}^r(m).$$

These events are all disjoint so we have $\mathbb{P}\left[\bigcup_{i=1}^5 B_i\right] = \sum_{i=1}^5 \mathbb{P}[B_i]$. Now,

$$\mathbb{P}[B_2] = \mathbb{P}[|c_{t_a}| = 2 | |e_{t_a}| \geq 3, l = 2] \mathbb{P}[|e_{t_a}| \geq 3, l = 2] \leq \frac{2}{3} \mathbb{P}[|e_{t_a}| \geq 3, l = 2],$$

by part 2 of Assumption 1.1. Next,

$$\mathbb{P}[B_3]$$

$$\begin{aligned} &= \mathbb{P}[a \in \{a_{n-1}, a_n\} | |e_{t_a}| \geq 3, |c_{t_a}| = 2, l \neq 2, l \notin 4\mathbb{N}] \mathbb{P}[|e_{t_a}| \geq 3, |c_{t_a}| = 2, l \neq 2, l \notin 4\mathbb{N}] \\ &\leq \frac{1}{3} \mathbb{P}[|e_{t_a}| \geq 3, |c_{t_a}| = 2, l \neq 2, l \notin 4\mathbb{N}], \end{aligned}$$

again by part 2 of Assumption 1.1 and noting that the first probability is maximised when $l = 6$. To deal with event B_4 , we condition on the values of the two sets $\{a_1, \dots, a_l\}$ and $\{a_1, \dots, a_{4j-4}\}$. Now a_{4j-3} is, by construction, the smallest element in $\{a_1, \dots, a_l\} \setminus \{a_1, \dots, a_{4j-4}\}$ and so can be identified under the conditioning, and a_{4j-2} is uniform on $\{a_1, \dots, a_l\} \setminus \{a_1, \dots, a_{4j-4}, a_{4j-3}\}$. Since a_{4j-1} is the smallest element in $\{a_1, \dots, a_l\} \setminus \{a_1, \dots, a_{4j-2}\}$, we see that, under this conditioning, the event that $a_{4j-1} > a_{4j-2}$ is the same as the event that a_{4j-2} is chosen to be the smallest element in $\{a_1, \dots, a_l\} \setminus \{a_1, \dots, a_{4j-3}\}$. Under the conditioning, this event has probability at most $1/3$ which is achieved when $l = 4$ (and so $j = 1$). We deduce that

$$\mathbb{P}[B_4] \leq \frac{1}{3} \mathbb{P}[|e_{t_a}| \geq 3, |c_{t_a}| = 2, \exists j \in \{1, \dots, \lfloor l/4 \rfloor\} : a \in \{a_{4j-3}, a_{4j-2}, a_{4j-1}, a_{4j}\}].$$

For the final event we have

$$\mathbb{P}[B_5] \leq \frac{3}{7} \mathbb{P}[|c_{t_a}| \geq 4],$$

since the worst possible case is when $|c_{t_a}| = 7$ (and so $d' = 1$). Combining these gives

$$\begin{aligned} & \mathbb{P}[F_a = * | \phi_a \neq \infty] \\ & \leq \mathbb{P}[|c_{t_a}| = 1] \\ & + \frac{2}{3} \left(\mathbb{P}[|e_{t_a}| \geq 3, l = 2] + \mathbb{P}[|e_{t_a}| \geq 3, |c_{t_a}| = 2, l \neq 2, l \notin 4\mathbb{N}] \right. \\ & \quad + \mathbb{P}[|e_{t_a}| \geq 3, |c_{t_a}| = 2, \exists j \in \{1, \dots, \lfloor l/4 \rfloor\} : a \in \{a_{4j-3}, a_{4j-2}, a_{4j-1}, a_{4j}\}] \\ & \quad \left. + \mathbb{P}[|c_{t_a}| \geq 4] \right) \\ & \leq \mathbb{P}[|c_{t_a}| = 1] + \frac{2}{3} (1 - \mathbb{P}[|c_{t_a}| = 1]) \leq \frac{11}{15}, \end{aligned}$$

by part 3 of Assumption 1.1. □

Proof of Proposition 7.8. We prove just the third statement, as the other two are similar. Let A_t^{EX} be a realisation of $\text{EX}(4, f, G)$ with $A_0^{\text{EX}} = \{a, b, c, d\}$. Then

$$\begin{aligned} & \mathbb{P}[|\{a, b, c, d\} \cap I(R)| = 1, |\{a, b, c, d\} \cap I(W)| = 3] \\ & = \mathbb{P}[|A_T^{\text{EX}} \cap R| = 1, |A_T^{\text{EX}} \cap W| = 3] \\ & \geq (1 - 2^{-9})^2 \frac{|R| \binom{|W|}{3}}{\binom{|V|}{4}}, \end{aligned}$$

where the inequality follows from Propositions 2.2 and 2.4, since $T = 20T_{\text{EX}(4,f,G)}$. □

References

- [1] David Aldous and James Allen Fill, *Reversible Markov Chains and Random Walks on Graphs* (2002), Unfinished monograph, recompiled 2014, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [2] N Berestycki and B Sengul, *Cutoff for conjugacy-invariant random walks on the permutation group*, *Probab. Theor. Rel. Fields* **173** (2019), no. 3–4, 1197–1241. MR-3936154
- [3] Pietro Caputo, Thomas M. Liggett, and Thomas Richthammer, *Proof of Aldous’ spectral gap conjecture*, *J. Amer. Math. Soc.* **23** (2010), no. 3, 831–851. MR-2629990
- [4] Fan Chung and Alexander Tsias, *Hypergraph coloring games and voter models*, *Internet Math.* **10** (2014), no. 1–2, 66–86. MR-3274540
- [5] Persi Diaconis and Mehrdad Shahshahani, *Generating a random permutation with random transpositions*, *Z. Wahrsch. Verw. Gebiete* **57** (1981), no. 2, 159–179. MR-0626813
- [6] L. Flatto, A. M. Odlyzko, and D. B. Wales, *Random shuffles and group representations*, *Ann. Probab.* **13** (1985), no. 1, 154–178. MR-0770635
- [7] Bob Hough, *The random k cycle walk on the symmetric group*, *Probab. Theor. Rel. Fields* **165** (2016), no. 1–2, 447–482. MR-3500276
- [8] Gabriel Istrate, Cosmin Bonchis, and Mircea Marin, *Interactive particle systems and random walks on hypergraphs* (2014). MR-3557646
- [9] Hubert Lacoïn, *Mixing time and cutoff for the adjacent transposition shuffle and the simple exclusion*, *Ann. Probab.* **44** (2016), no. 2, 1426–1487. MR-3474475
- [10] David A Levin, Yuval Peres, and Elizabeth L Wilmer, *Markov Chains and Mixing Times*, American Mathematical Society, Providence, RI (2009). MR-2466937
- [11] Thomas M. Liggett, *Interacting particle systems*, *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 276, Springer-Verlag, New York, 1985. MR-0776231
- [12] Thomas M. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes*, *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 324, Springer-Verlag, Berlin, 1999. MR-1717346

- [13] Nathan Lulov and Igor Pak, *Rapidly mixing random walks and bounds on characters of the symmetric group*, J. Algebraic Combin. **16** (2002), no. 2, 151–163. MR-1943586
- [14] Ben Morris, *The mixing time for simple exclusion*, Ann. App. Probab. **16** (2006), no. 2, 615–635. MR-2244427
- [15] Roberto Imbuzeiro Oliveira, *Mixing of the symmetric exclusion processes in terms of the corresponding single-particle random walk*, Ann. Probab. **41** (2013), no. 2, 871–913. MR-3077529